arXiv:1403.6532v1 [math.ST] 25 Mar 2014

Minimax Optimal Rates for Poisson Inverse Problems with Physical Constraints

Xin Jiang[1] Garvesh Raskutti[2] Rebecca Willett [3]

[1] Department of Electrical and Computer Engineering, Duke University
[2] Department of Statistics, University of Wisconsin-Madison
[3] Department of Electrical and Computer Engineering, University of Wisconsin-Madison

Abstract

This paper considers fundamental limits for solving sparse inverse problems in the presence of Poisson noise with physical constraints. Such problems arise in a variety of applications, including photon-limited imaging systems based on compressed sensing. Most prior theoretical results in compressed sensing and related inverse problems apply to idealized settings where the noise is i.i.d., and do not account for signal-dependent noise and physical sensing constraints. Prior results on Poisson compressed sensing with signal-dependent noise and physical constraints in [1] provided upper bounds on mean squared error performance for a specific class of estimators. However, it was unknown whether those bounds were tight or if other estimators could achieve significantly better performance. This work provides minimax lower bounds on mean-squared error for sparse Poisson inverse problems under physical constraints. Our lower bounds are complemented by minimax upper bounds. Our upper and lower bounds reveal that due to the interplay between the Poisson noise model, the sparsity constraint and the physical constraints: (i) the mean-squared error does not depend on the sample size n other than to ensure the sensing matrix satisfies RIP-like conditions and the intensity T of the input signal plays a critical role; and (ii) the mean-squared error has two distinct regimes, a low-intensity and a high-intensity regime and the transition point from the low-intensity to high-intensity regime depends on the input signal f^*. In the low-intensity regime the mean-squared error is independent of T while in the high-intensity regime, the mean-squared error scales as $\frac{s \log p}{T}$, where s is the sparsity level, p is the number of pixels or parameters and T is the signal intensity.

1 Introduction

In this paper we investigate minimax rates associated with Poisson inverse problems under physical constraints, with a particular focus on compressed sensing (CS) [2, 3]. The idea behind compressed sensing is that when the underlying signal is sparse in some basis, the signal can be recovered with relatively few sufficiently diverse projections. While current theoretical CS performance bounds are very promising, they are often based on assumptions such as additive signal-independent noise which do not hold in many realistic Poisson inverse problem settings.

Poisson inverse problems arise in a variety of applications.

- Several imaging systems based on compressed sensing have been proposed. The Rice single-pixel camera [4] or structured illumination fluorescence microscopes [5] collect projections of a scene sequentially, and collecting a large number of observations within a limited time frame necessarily limits the number of photons that are detected and hence the signal-to-noise ratio (SNR). In particular, if we collect n projection measurements of a scene with p pixels over a span of T seconds, then the average number of photon counts per measurement scales like T/n. Thus even if the hardware used to determine which projections are measured (*e.g.,* a digital micromirror device) is capable of operating at very high speeds, we will still face photon limitations and hence Poisson noise [6].

- In conventional fluorescence microscopy settings (*i.e.,* non-compressive, direct image acquisition) we may wish to represent a signal using a sparse superposition of Fourier and pixel elements to facility quantitative tissue analysis [7, 8, 9]. In this context one may wish to characterize the reliability of such

a decomposition to ensure sufficient acquisition time or sufficient quantities of fluorescent protein. In this context n is the number of pixels observed, and p is the number of coefficients to be estimated.

- In network flow analysis, we wish to reconstruct average packet arrival rates and instantaneous packet counts for a given number of streams (or flows) at a router in a communication network, where the arrivals of packets in each flow are assumed to follow a Poisson process. All packet counting must be done in hardware at the router, and any hardware implementation must strike a delicate balance between speed, accuracy, and cost [10, 11, 12, 13]. In this context, n is the number of (expensive) memory banks in a router and ideally will be much less than p, the total number of flows to be estimated. T corresponds to the total number of packets observed.

- Poisson inverse problems also arise in DNA analysis [14], pediatric computer-aided tomography [15], Twitter data analysis [16], and quantitative criminology [17].

In all the above settings, when small numbers of events (*e.g.,* photons, x-rays, or packets) are observed, the noise or measurement error can be accurately modeled using a Poisson distribution [6], so recovering a signal from projection measurements amounts to solving a Poisson inverse problem. Furthermore, the Poisson inverse problem is subject to physical constraints (e.g. hardware constraints, flow constraints) which are outlined below. In general, the available data acquisition time places fundamental limits on our ability to accurately reconstruct a signal in a Poisson noise setting. Until now, relatively little about the nature of these limits was understood.

1.1 Observation model and physcial constraints

In this paper we consider an observation model of the form

$$y \sim \text{Poisson}(TAf^*), \tag{1}$$

where $A \in \mathbb{R}_+^{n\times p}$ in a sensing matrix corresponding to the n different projections of our signal of interest $f^* \in \mathbb{R}_+^p$ and $T \in \mathbb{R}_+$ is the total data acquisition time. In particular, (1) is a shorthand expression for the model

$$y_i \sim \text{Poisson}\left(T\sum_{j=1}^{p} A_{i,j}f_j^*\right), \qquad i \in \{1,\dots,n\},$$

where the y_i's are independent. Since for the Poisson distribution the variance depends on the signal intensity, the noise of the observed signal y depends on T, A and f^*.

In particular, f^* corresponds to the rate at which events are being generated, which are intrinsically nonnegative. In addition, A must be composed of nonnegative real numbers with each column summing to at most one. Specifically, A must satisfy the following physical constraints:

$$A \succeq 0 \qquad \text{nonnegative;} \tag{2a}$$

$$\mathbb{1}_{n\times 1}^\top A \preceq \mathbb{1}_{p\times 1} \qquad \text{bounded column sums.} \tag{2b}$$

Together these constraints imply (a) $A_{i,j} \in [0,1]\ \forall(i,j)$ and (b) $\|Af\|_1 \le \|f\|_1$ for all nonnegative f; this latter property with equality is often called *flux preservation*, particularly in the astronomy community. Thoughtout this paper, we also refer to constraint (2b) as the flux-preserving constraint.

To build intuition about these constraints it is helpful to consider an imaging system in which $A_{i,j}$ corresponds to the likelihood of a photon generated at location j at the source f^* hitting our detector at location i. Such a likelihood is necessarily nonnegative and bounded, and the j^{th} column sum, which corresponds to the likelihood of a photon from location j hitting any of the n detectors, must be bounded by one. Without the constraint (2b), we could have an unrealistic imaging system in which more photons

are detected than we have incident upon the aperture. These constraints arise naturally in all the motivating applications described above using similar arguments.

We assume $\|f^*\|_1 = 1$ throughout, so the total expected number of observed events is proportional to T; thus T alone (*i.e.,* neither f^* nor A) controls the signal-to-noise ratio. This model is equivalent to assuming the signal of interest integrates to one (like a probability density), and by allotting ourselves T units of time to acquire data (*e.g.,* collect photons).

To incorporate sparsity, we assume f^* is $(s+1)$-*sparse* in a basis spanned by the columns of the orthonormal matrix $D \in \mathbb{R}^{p\times p}$, where we write $D = [d_1, \ldots, d_p]$ for $d_i \in \mathbb{R}^p$ for all i. We assume that $d_1 = p^{-1/2}\mathbb{1}_{p\times 1}$. [1] Standard choice of basis matrices D are Fourier and wavelet basis matrices which we discuss in more detail later. For an arbitrary vector $v \in \mathbb{R}^p$, let $\|v\|_0 \triangleq \sum_{j=1}^p \mathbf{1}(v_j \neq 0)$ denote the standard ℓ_0-norm.

Hence, if we define $\theta^* \triangleq D^\top f^* \in \mathbb{R}^p$, then $\|\theta^*\|_0 \leq s+1$. Note that by construction, $\theta^*_1 \equiv \langle d_1, f^*\rangle \equiv 1/\sqrt{p}$. Let $\bar{D} \triangleq [d_2, d_3, \ldots, d_p] \in \mathbb{R}^{p\times(p-1)}$ be all but the first (constant) basis vector, let $\bar{\theta}^* \triangleq \bar{D}^\top f^*$ be all but the first (known) coefficients, and note that $\bar{\theta}^*$ is s-sparse.

To summarize, we consider f^* belonging to the set:

$$\mathcal{F}_{p,s,D} \triangleq \left\{ f \in \mathbb{R}^p_+ \;:\; \|f\|_1 = 1,\; \|\bar{D}^\top f\|_0 \leq s \right\}.$$

We assume throughout the paper that $f^* \in \mathcal{F}_{p,s,D}$.

1.2 Our Contributions

In this paper, we derive mean-squared error bounds for the Poisson inverse problem described in (1) where the function f^* belongs to $\mathcal{F}_{p,s,D}$. Our main results are presented in Section 2. In particular, in Section 2.2 we provide a minimax lower bound for the mean-squared error, restricted to the set of signals and sensing matrices that satisfy the above physical constraints. In Section 2.3 we provide a minimax upper bound based on analysis of a modified version of the estimator considered in [1].

Our upper and lower bounds reveal two new effects due to the interplay between the Poisson model, the sparsity assumption and the physical constraints. The first new effect is that the flux-preserving constraint ensures that the mean-squared error has no dependence on the sample size n outside of needing n to be sufficiently large for A to satisfy RIP-like assumptions. Instead the intensity T plays a significant role in the mean-squared error. Secondly, the interplay between the non-negativity constraint and the sparsity assumption means that there are two distinct regimes, a *low-intensity* and *high-intensity* regime. In the low-intensity regime, *i.e.,* for T below some critical level, the mean-squared error is independent of T, whereas in the high-intensity regime the mean-squared error scales as $\frac{s\log p}{T}$, which we prove is the minimax optimal rate. The point at which the mean-squared error transitions from the low-intensity to the high-intensity regime depends significantly on the input signal f^* and the orthogonal basis matrix that induces sparsity on f^*. Hence our upper and lower bounds are sharp up to constant in the high-intensity regime whereas in the low-intensity regime, our upper and lower bounds do not match since the optimal choices of f^* depend significantly on the orthonormal basis matrix D. The theoretical performance bounds are supported by a suite of simulations presented in Section 3. Thus the main theoretical results in this paper are perhaps surprising and certainly at odds with similar results appearing in Gaussian or bounded noise settings where the bounds are independent of the sparsifying basis and the choice of f^*, and the mean-squared error bounds depend on n.

2 Assumptions and Main Results

We provide upper and lower bounds for the minimax mean-squared error. Let $\widehat{f} \equiv \widehat{f}(y)$ denote an estimator of f^* (*i.e.,* $\widehat{f}$ is a measurable function of y, and TAD). The performance of the estimator is evaluated using

[1] This assumption precludes signals that are sparse in the canonical or pixel basis. Our proof techniques, with some technical adjustments, can be adapted to that setting.

the following squared ℓ_2 risk:

$$\mathcal{R}(\widehat{f}, f^*) \triangleq \mathbb{E}[\|\widehat{f} - f^*\|_2^2],$$

where the expectation is taken over y. The minimax risk is defined as:

$$\min_{\widehat{f}} \max_{f^* \in \mathcal{F}_{p,s,D}} \mathcal{R}(\widehat{f}, f^*).$$

Here the $\min$ is taken over the set of measurable functions of (TAD, y). To begin we state assumptions imposed on the sensing matrix A.

2.1 Assumptions on A

Assumption 2.1. *There exist constants a_ℓ and a_u such that $a_\ell < a_u$ and a matrix $\widetilde{A} \in [a_\ell/\sqrt{n}, a_u/\sqrt{n}]^{n\times p}$*

$$A = \frac{\widetilde{A} + \frac{a_u - 2a_\ell}{\sqrt{n}} \mathbb{1}_{n\times p}}{2(a_u - a_\ell)\sqrt{n}}. \tag{3}$$

Assumption 2.2. *For all $u \in \mathbb{R}^p$, $\|u\|_0 \le 2s$, there exists a constant $\delta_s \equiv \delta_s(n,p) > 0$ such that*

$$\|\widetilde{A}Du\|_2^2 \le (1+\delta_s)\|u\|_2^2.$$

Assumption 2.3. *For all $u \in \mathbb{R}^p$, $\|u\|_0 \le 2s$, there exists a constant $\delta_s \equiv \delta_s(n,p) > 0$ such that*

$$\|\widetilde{A}Du\|_2^2 \ge (1-\delta_s)\|u\|_2^2.$$

Assumption 2.1 ensures that $A \succeq 0$. Note that in many practical cases in which Assumptions 2.1, 2.2, and 2.3 are satisfied, we will have $a_\ell < 0 < a_u$. Assumptions 2.2 and 2.3 together are the the RIP conditions for $\widetilde{A}D$ introduced in [18]. Our minimax lower bound requires only Assumption 2.2 whereas our upper bound requires both Assumption 2.2 and Assumption 2.3, which is consistent with [19].

Assumption 2.1 also introduces a new matrix $\widetilde{A}$. This reflects the actual process of constructing a sensing matrix that satisfies the Conditions (2a) and (2b). The typical design process would be to start with a bounded matrix which satisfies the RIP conditions, rescale each entry and add an offset to the original matrix to make sure the new matrix is positive and flux-preserving. The following lemma ensures that these two constraints are satisfied via such operations:

Lemma 2.1. *If the matrix $\widetilde{A}$ satisfies $\widetilde{A}_{i,j} \in \frac{1}{\sqrt{n}}[a_\ell, a_u]$ for all entries, and the sensing matrix A is generated according to Eq. (3), then A satisfies Conditions (2a) and (2b).*

The proof is provided in Section A.1.

For example, let A be constructed based on a shifted and rescaled Bernoulli ensemble matrix $\widetilde{A}$, where

$$\mathbb{P}(\widetilde{A}_{i,j}) = \begin{cases} 1/2, & \widetilde{A}_{i,j} = -\sqrt{\frac{1}{n}} \\ 1/2, & \widetilde{A}_{i,j} = \sqrt{\frac{1}{n}}. \end{cases} \tag{4}$$

The resulting sensing matrix A has i.i.d. entries valued $1/(2n)$ or $1/n$ and using results in [20], $\widetilde{A}$ and $\widetilde{A}D$ satisfy the RIP conditions in Assumptions 2.2 and 2.3 for $n \ge C_0 s \log(p/s)$ with probability at least $1 - e^{-C_1 n}$.

2.2 Lower bounds on minimax risk

In this section we present a lower bound on the minimax risk. The interaction between the orthonormal basis matrix D and the sparsity constraint has an effect on the lower bound which is captured by the following s-sparse localization quantity:

Definition 2.1 (s-sparse localization)**.** *λ_s is said to be the s-sparse localization quantity of a matrix X if*

$$\lambda_s = \lambda_s(X) \triangleq \max_{\substack{v \in \{-1,0,1\}^p \\ \|v\|_0 = s}} \|Xv\|_\infty.$$

The name "localization" in Definition 2.1 derives from a similar quantity defined in [21]. The key difference between our localization constant and the one in [21] is that we incorporate the sparsity level directly within the definition. Our minimax lower bound depends on $\lambda_k(\bar{D})$ for $1 \le k \le s$.

Theorem 2.1 (Minimax lower bound)**.** *For $k = 1, \dots, s$, let $\lambda_k = \lambda_k(\bar{D})$ be the k-sparse localization quantity of $\bar{D}$. If $p \ge 10$, $1 \le s < p/3 - 1$, and Assumptions 2.1 and 2.2 hold with $0 \le \delta_s < 1$, then there exists a constant $C_L > 0$ that depends only on a_u and a_ℓ such that*

$$\min_{\widehat{f}} \max_{f^* \in \mathcal{F}_{p,s,D}} \mathcal{R}(\widehat{f}, f^*) \ge C_L \max_{1 \le k \le s} \left\{ \min \left(\frac{k}{p^2 \lambda_k^2}, \frac{k}{(1+\delta_s)T} \log \frac{p-k-1}{k/2} \right) \right\}. \tag{5}$$

The proof is provided in Section 4.1 and involves adapting the information-theoretic techniques developed in [19, 22, 23] to the Poisson noise setting.

Remark 2.1. *Within the* max *term, the lower bound is a minimum of two terms. These two terms represent the minimax rate in a "low-intensity" regime where T is less than $O(\min_{k \le s} \lambda_k^2 p^2 \log p)$ and a "high-intensity" regime where T is above $O(\max_{k \le s} \lambda_k^2 p^2 \log p)$. The sparse localization quantity λ_k plays a significant role in the performance at lower intensities, whereas at high intensities the scaling is simply $O(\frac{s \log p}{T})$.*

Remark 2.2. *It should be noted that the quantity λ_k is not a constant in general. It is a function of D and k, which also means that λ_k can be a function of the signal dimension p. Depending on the choice of the basis, λ_k can be on the order of $k/\sqrt{p}$ (*e.g.*, for the DCT basis), or a constant (*e.g.*, for the Haar wavelet basis). The computation of λ_k for various D is described in Section 3.1. As a result, our lower bound in the low-intensity setting for the discrete cosine basis is different from the Haar wavelet basis (see Table 1).*

Remark 2.3. *The bound in (5) is independent of n, but note that the assumptions only hold for n sufficiently large. For instance, if A were generated using a Bernoulli ensemble as in (4), we would need $n = O(s \log(p/s))$ to ensure Assumptions 2.2 and 2.3 hold.*

2.3 Upper bound

Now we provide an upper bound on the minimax rate. To obtain the upper bound, we analyze an ℓ_0-penalized likelihood estimator $\widehat{f}$ introduced in [1].

Theorem 2.2 (Minimax upper bound)**.** *If $f^* \in \mathcal{F}_{p,s,D}$, A satisfies Assumption 2.1, 2.2, and 2.3 with $0 \le \delta_s < 1$, then there exists a constant $C_U > 0$ depending only on a_u and a_ℓ such that*

$$\min_{\widehat{f}} \max_{f^* \in \mathcal{F}_{p,s,D}} \mathcal{R}(\widehat{f}, f^*) \le C_U \min \left(\frac{s \log_2 p}{(1-\delta_s)T} + \frac{s \log_2(T+1)}{(1-\delta_s)T}, s \max_{j,k} |D_{j,k}|^2, 1 \right).$$

Additionally, as long as T does not increase exponentially with p, $\log_2(T+1)$ is dominated by $c\log_2 p$ for some $c > 0$, our upper bound satisfies

$$\min_{\widehat{f}} \max_{f^* \in \mathcal{F}_{p,s,D}} \mathcal{R}(\widehat{f}, f^*) \leq C'_U \min\left(\frac{s\log_2 p}{(1-\delta_s)T}, s\max_{j,k}|D_{j,k}|^2, 1\right). \tag{6}$$

The proof is provided in Section 4.2.

Remark 2.4. *The first term $\frac{s\log p}{T}$ arises from analyzing a penalized maximum likelihood estimator. The second and third term arises from using the 0 estimator and upper bounding $\|f^*\|_2^2$ by $\min(s\max_{j,k}|D_{j,k}|^2, 1)$.*

Remark 2.5. *Like the lower bound, we see different behavior of the reconstruction error depending on the intensity T. For $T > \min(s\log p, (\max_{i,j}|D_{j,k}|)^{-2}\log p)$ the upper bound scales as $\frac{s\log p}{T}$ while for $T \leq \min(s\log p, (\max_{j,k}|D_{j,k}|)^{-2}\log p)$ the upper bound is independent of T. Note that in the high-intensity regime (large T), the upper and lower bounds of $\frac{s\log p}{T}$ match up to a constant. We discuss the scaling in the low-intensity regime in more detail in Section 3.1.*

Remark 2.6. *The lower bound in Theorem 2.1 only requires Assumptions 2.1 and 2.2, whereas the upper bound requires the additional Assumption 2.3. This is consistent with assumptions required for the minimax upper and lower bounds in [19].*

Remark 2.7. *Like in the lower bound, the number of observations n does not appear in the mean-squared error expression.*

3 Consequences and discussion

The minimax lower bounds show that the mean-squared error scales as $\frac{s\log(p/s)}{T}$ provided T is sufficiently high and the upper bound scales as $\frac{s\log p}{T}$, where T corresponds to the data acquisition time or the total number of event available to sense. Notice that the number of observations n does not appear in the mean-squared error expression. Hence increasing the number of observations n beyond what is necessary to satisfy the necessary assumptions on A is not going to reduce the mean-squared error. Furthermore, our analysis proves that the mean-squared error is inversely proportional to the intensity T, so T plays the role that n usually does in the standard compressed sensing setup. We also carefully characterize the mean-squared error behavior when the intensity T is low, and show *bounds that depend on the interactions between the sparsifying basis and the physical constraints.* We elaborate on these concepts below.

3.1 High and low intensities

As mentioned earlier, the lower bounds in Theorem 2.1 and the upper bounds in Theorem 2.2 agree up to a constant when the intensity is sufficiently high. In this section, we explore the relationship between the upper and lower bounds in the low-intensity setting for discrete cosine transform (DCT), discrete Hadamard transform (DHT) and discrete Haar wavelet basis (DWT). The elements for the DCT basis matrix are:

$$D^{\mathrm{DCT}}_{j,k} = \begin{cases} \sqrt{\frac{1}{p}}, & k = 1, j = 1, \ldots, p \\ \sqrt{\frac{2}{p}}\cos\left(\frac{(2j-1)(k-1)\pi}{2p}\right), & k = 2, \ldots, p, j = 1, \ldots, p. \end{cases} \tag{7}$$

The DHT matrix $D^{\mathrm{DHT}}_{(m)} \in \mathbb{R}^{2^m \times 2^m}$ has the form

$$\left(D^{\mathrm{DHT}}_{(m)}\right)_{j,k} = \frac{1}{2^{\frac{m}{2}}}(-1)^{j\cdot k}. \tag{8}$$

The DWT matrix of dimension 2^m by 2^m, the non-zeros have magnitudes in the set $\{2^{-m/2}, 2^{-(m-1)/2}, \ldots, 2^{-1/2}\}$.

In the low-intensity setting, the lower bound scales like $\max_{1\le k\le s} \frac{k}{p^2\lambda_k^2}$, while the upper bound scales as $\min\left(1, s\max_{j,k}|D_{j,k}|^2\right)$. Since the bounds depend on λ_k and $\max_{j,k}|D_{j,k}|$, it is not straightforward to see how these two bound match with each other. In Table 1, we show the comparison of the bounds for DCT basis, DHT basis and DWT basis. The calculation of λ_k can be found in Section B.2.

	λ_k	Lower bound	$\max_{j,k}\lvert D_{j,k}\rvert$	Upper bound
DCT & DHT	$\frac{\sqrt{2}k}{\sqrt{p}}$	$\min\left(\frac{1}{p}, \frac{s\log(p/s)}{T}\right)$	$\frac{2}{\sqrt{p}}$	$\min\left(\frac{s}{p}, \frac{s\log p}{T}\right)$
DWT	$\frac{1}{\sqrt{2}-1}$	$\min\left(\frac{s}{p^2}, \frac{s\log(p/s)}{T}\right)$	$\frac{1}{\sqrt{2}}$	$\min\left(1, \frac{s\log p}{T}\right)$

Table 1: Lower and upper bounds for three different bases (neglecting constants).

For the high-intensity setting, the lower and upper bound have the same rate regardless of the basis (except for the minor difference in the log term), which proves the bound is tight. For the low-intensity setting, the upper and lower bounds do not match.

To see why this is so, note that different sparse supports for θ yield very different MSE performances because of the interaction between the sparse support and the nonnegativity constraints. Our proof shows that in low-intensity settings, the MSE is proportional to the squared ℓ_2 norm of the zero-mean signal $\|f^* - \mathbb{1}_{n\times 1}/\sqrt{p}\|_2^2 = \|\bar{\theta}^*\|_2^2$. Upper and lower bounds on this quantity depend on the interaction between the constraint, $\|f^*\|_1 = 1$, the sparsity constraint, and the orthonormal basis matrix D. The general lower bound $\frac{k}{p^2\lambda_k^2}$ derives from the following θ^*:

$$\theta_j^* = \begin{cases} \frac{1}{\sqrt{p}}, & j = 1 \\ \pm\frac{1}{p\lambda_k}, & 2 \le j \le k+1 \\ 0, & \text{otherwise.} \end{cases} \tag{9}$$

It is straightforward to see that $\|\bar{\theta}^*\|_2^2 = \frac{k}{p^2\lambda_k^2}$ and Section 4.1 shows that $f^* = D\theta^*$ satisfies all the necessary constraints. On the other hand, the upper bound $\max(1, s\max_{j,k}|D_{j,k}|^2)$ follows from

$$\|\bar{\theta}^*\|_2^2 < \|\theta^*\|_2^2 = \|f^*\|_2^2 \le \|f^*\|_1^2 = 1$$

and

$$\|\bar{\theta}^*\|_2^2 \le s\|\theta^*\|_\infty^2 = s\|D^T f^*\|_\infty^2 \le s\max_{j,k}|D_{j,k}|^2.$$

Whether the upper or lower bounds are tight depend on the orthonormal basis matrix D. For example for the DWT basis, when s is sufficiently large (*i.e.,* $\Omega(\log p)$), the ℓ_2 norm of $f \in \mathcal{F}_{p,s,D}$ can be large, leading to tighter lower bounds.

Lemma 3.1. *Let D be the discrete Haar wavelet basis matrix. If $D^\top f^*$ is $(s+1)$-sparse with $p \ge 4$, $s \ge \log_2 p$, $T \le \frac{(a_u - a_\ell)^2 \log p}{4(1+\delta_s)}$, and Assumptions 2.1 and 2.2 hold with $0 \le \delta_s < 1$, then there exists an absolute constant $C_L'' \ge 1/8$ such that*

$$\min_{\widehat{f}} \max_{f^* \in \mathcal{F}_{p,s,D}} \mathbb{E}\left[\|f^* - \widehat{f}\|_2^2\right] \ge C_L''. \tag{10}$$

The proof is provided in Appendix A.2.

To provide an explanation for the effect of the matrix D on the minimax rates, the nonnegativity constraints implicitly impose limits on the amplitudes of the coefficients of f^*, and these limits strongly impact

MSE performance. For instance, for 1-sparse $\bar{\theta}$ in the DWT basis, fine-scale basis vectors have amplitude of $O(1)$, and so the corresponding basis coefficient cannot be too large without violating the nonnegativity constraint; in contrast, coarse-scale basis vectors have amplitude of $O(1/\sqrt{p})$ and can have much larger coefficients without violating the nonnegativity constraint. For the DWT basis, a sparse support can be chosen to generate a high-ℓ_2 norm signal that approximates a delta function in $\mathcal{F}_{p,s,D}$; this kind of signal dominates our upper bounds. A different sparse support set may correspond to basis vectors with mostly disjoint supports; such a signal must have a small ℓ_2 norm in order to be in $\mathcal{F}_{p,s,D}$ (i.e. to be nonnegative); this kind of signal dominates our lower bounds. A similar effect occurs in the DCT basis. We may contrast these behaviors with what happens in a conventional Gaussian noise analysis without a nonnegativity constraint, where we would have no constraints on coefficient magnitudes to ensure nonnegativity. This example and a similar example for the DCT basis are presented with a simple simulation described in Section 3.2.1.

3.2 Simulation results

We further assess the proposed bounds by a series of experiments that examine the reconstruction performance of signals that satisfy our assumptions. We assume the signal is s-sparse under the DCT or DWT basis. Plots are not generated for the DHT, as it exhibits similar performance behavior to the DCT.

We construct the signal θ^* as follows: the first coefficient $\theta_1^* = 1/\sqrt{p}$ corresponds to the DC level and ensures $\|D\theta^*\|_1 = 1$. The locations of the remaining s non-zero coefficients are randomly generated, while the amplitude of each non-zero coefficient is the same and chosen to be $\pm 1/(p\lambda_s)$, which ensures that $D\theta^* \in \mathcal{F}_{p,s,D}$. (In particular, this construction of θ^* corresponds to elements of the packing set described in Section 4.1.) The sensing matrix A is generated according to Eqs. (3) and (4). The reconstruction is done by solving a ℓ_1-penalized Poisson likelihood optimization problem. The reason for using a relaxed ℓ_1 penalty instead of the ℓ_0 penalty which is used for the analysis of the upper bound is that the ℓ_0 penalization function is non-convex, and cannot be solved efficiently. The optimization problem is solved using the SPIRAL algorithm developed in [24]. Each data point in the plot shows the mean-squared error (MSE) averaged over 100 experiments.

3.2.1 Impact of sparsifying dictionary on rates

As mentioned above, the the nonnegativity constraints impose limits on the amplitudes of the coefficients of f^*, and these limits strongly impact MSE performance. We illustrate that effect here. In particular, we see that randomly selected sparse supports and uniform coefficient magnitudes result in MSE rates which correspond to our lower bounds. In contrast, carefully constructed sparse supports (specific to the underlying sparsifying dictionary) and highly non-uniform coefficient magnitudes result in slower MSE rates which correspond to our upper bounds.

Figure 1 shows MSE behavior in the low-intensity regime for both these signal types in the DCT and DWT bases. All the signals used in the simulation belong to the function class $\mathcal{F}_{p,s,D}$, but the signals' sparse supports and energies vary. For Figures 1a, 1c and 1e the signals are sparse in the DCT basis, and for Figures 1b, 1d and 1f the signals are sparse in the DWT basis. In the plots, signals θ_1^* and θ_3^* are chosen to make $f_i^* = D\theta_i^*$ for $i = 1, 3$ approximate a δ-function (and thus have a large ℓ_2 norm). Signals θ_2^* and θ_4^* have randomly selected supports and non-zero coefficient magnitudes equal to $1/(p\lambda_s)$, as described above. None of these signals can be rescaled to have larger ℓ_2 norms without violating the nonnegativity constraints. Signals θ_2^* and θ_4^* control the rates in the lower bound, while signals θ_1^* and θ_3^* control the rates in the upper bound. Details on how the different signals were generated are in Appendix B.1.

3.2.2 MSE vs. T

As discussed before, the MSE's behavior changes under different intensities. In the high-intensity setting, theoretical bounds predict that MSE will be proportional to $1/T$. In the low-intensity setting, however, T is not a dominating factor in the lower bound – in fact, our bound is independent of T when T is below a

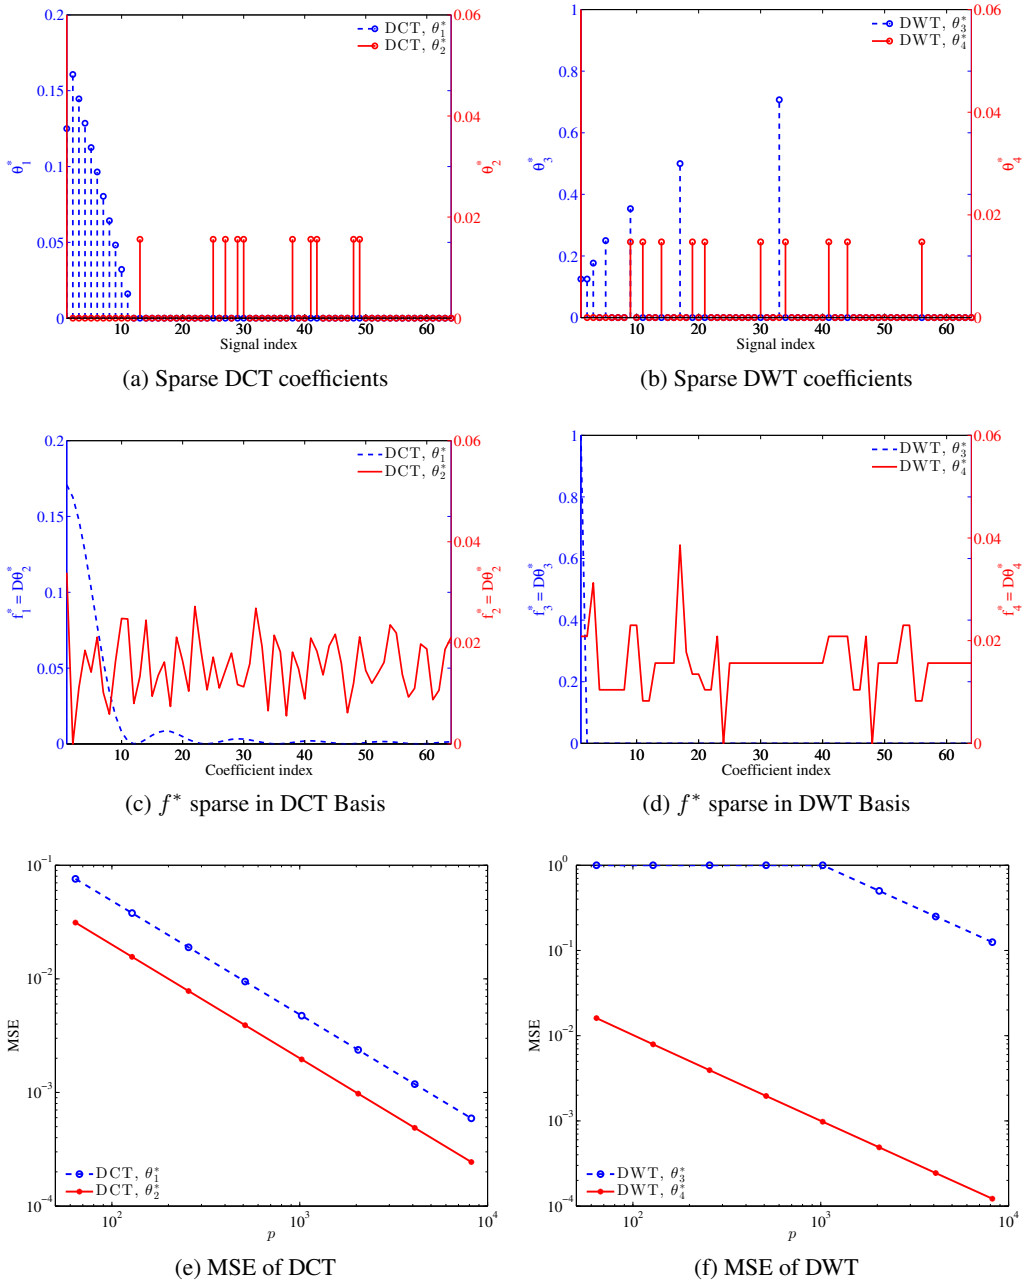

(a) Sparse DCT coefficients

(b) Sparse DWT coefficients

(c) f^* sparse in DCT Basis

(d) f^* sparse in DWT Basis

(e) MSE of DCT

(f) MSE of DWT

Figure 1: Signals and MSE scaling in the low-intensity regime with different sparse supports. θ_1^* and θ_3^* are chosen to make f^* approximate a δ-function (and thus have a large ℓ_2 norm), while sparse supports 2 and 4 are randomly chosen. None of these signals can be rescaled to have larger ℓ_2 norms without violating the nonnegativity constraints. $f_i^* = D\theta_i^*, i = 1, \ldots, 4$ are all in $\mathcal{F}_{p,s,D}$; that is, they are nonnegative with unit ℓ_1 norm. Signals like f_2^* and f_4^* control the rates in the lower bound, while signals like f_1^* and f_3^* control the rates in the upper bound. All the signals used in the experiment are 10-sparse.

critical, sparsifying-basis dependent threshold.

In Fig. 2, the plot shows how the total intensity affects the performance. "Elbows" can be observed in both lines for the DCT basis and the DWT basis, which indicate the behavior change predicted by our

theory. Note that in Fig. 2 we are plotting in log-log scale, and MSE $\propto \frac{1}{T} \Leftrightarrow \log(\text{MSE}) \propto -\log T$, which is exactly the linear relationship we observe when T is large. Once T drops below the critical value (left hand side of the plot), the MSE does not change for different T, which reflects the bound not depending on T in the low-intensity regime. The dashed lines are the value of $\|\bar{\theta}^*\|_2^2$ used in the experiments, which is an upper bound on the MSE. (The bound in Thm. 2.2 reflects an upper bound on $\|\bar{\theta}^*\|_2^2$.) Thus this plot suggests that MSE is proportional to $\|\bar{\theta}^*\|_2^2$ for small T as suggested by our discussion.

It should also be noted that the "elbow" for the DCT basis arrives at smaller T than for the DWT basis, as predicted recalling that the transition between low- and high-intensity regimes occurs when $T \propto \lambda^2(D)$, and $\lambda(\bar{D}^{\text{DWT}}) > \lambda(\bar{D}^{\text{DCT}})$ when $s < \sqrt{p}$.

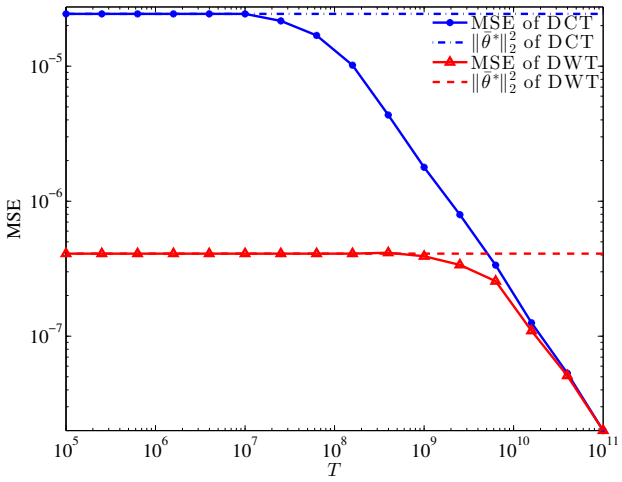

Figure 2: MSE vs. T. The plot shows MSE of signals sparse under the DCT basis and of signals sparse under the DWT basis. The dashed lines are the values of $\|\bar{\theta}^*\|_2^2$ used in the experiments; this value is also an upper bound of the MSE, and we see that this bound is tight at low intensities.

3.2.3 MSE vs. n

One practical question that arises in many Poisson inverse problems is how to best trade off between the number of measurements and the amount of time spent collecting each measurement (*i.e.,* the SNR of each measurement). Is it better to have a lot of noisy measurements, or to have a small number of high SNR measurements? The bounds derived in this paper help address this question.

In particular, Theorems 2.1 and 2.2 suggest that the upper and lower bounds are independent of n as long as n is large enough to ensure that the sensing matrix A satisfies Assumptions 2.2 and 2.3, an observation conjectured in [1]. This is contrary to most compressed sensing results, where the error usually goes to zero as the number of measurements increases. This interesting effect (or non-effect) of n is a direct result of the flux-preserving constraint, which enforces that as n increases, the number of events detected per sensor decreases. In other words, when T is held as a constant, the increase of n does not increase the overall signal-to-noise ratio. Instead, the energy is spread out onto more observations, which may bring more information about the signal, but also causes the noise-per-observation to rise.

This result suggests that once n is sufficiently large to ensure Assumptions 2.2 and 2.3 are satisfied, there is no advantage to increasing n further. In other words, a small number of high SNR measurement is sufficient, and often better in terms of hardware costs and reconstruction computational complexity. This result is similar to the finding recently reported in literature exploring the effects of quantization on compressed sensing measurements [25, 26].

These theoretical findings are illustrated in the below simulation. Fig. 3 displays MSE of the reconstruc-

tions as a function of n. The reconstruction error keeps almost constant as n varies within the error bars, which is consistent with the bounds where n has little influence.

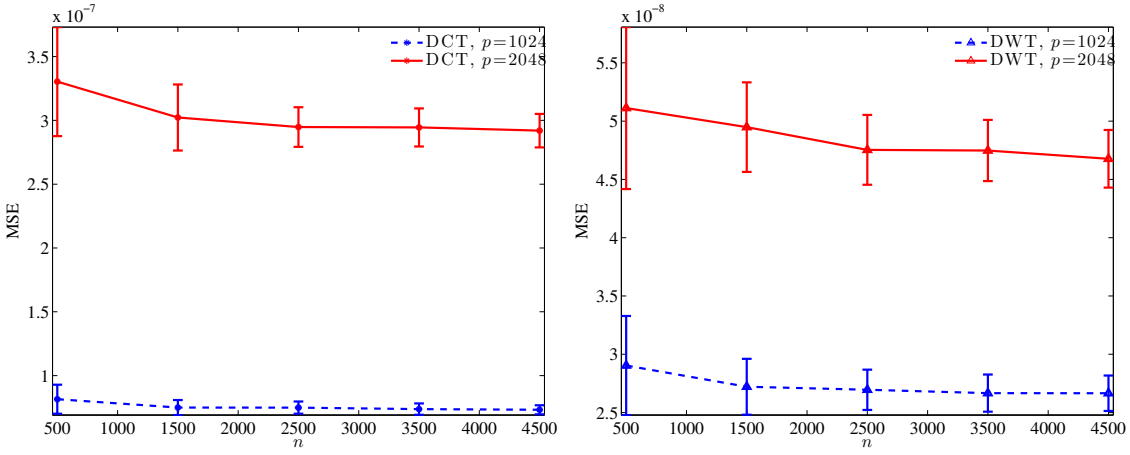

Figure 3: MSE vs. n ($T = 10^{10}$, $s = 10$). The reconstruction error keeps almost constant as n varies within the error bars, which is consistent with the bounds where n has little influence.

3.2.4 MSE vs. s

Our bounds in Thms. 2.1 and 2.2 suggest the MSE scales linearly with s at high intensities (*i.e.,* when T is large). In Fig. 4, MSE in the high-intensity ($T = 10^{12}$) is shown for both the DCT and DWT bases for $n = 500$. With both bases, the MSE grows linearly with s across a range of different p, which is consistent with the theoretical results.

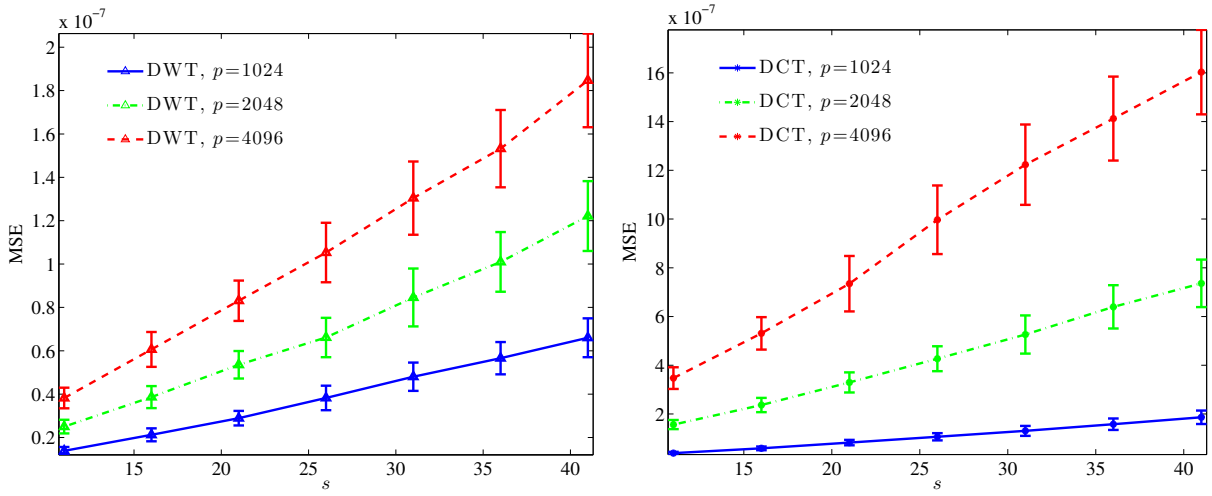

(a) DCT sparsity

(b) DWT sparsity

Figure 4: MSE vs. s in the high-intensity setting ($T = 10^{12}$ and $n = 500$). With both bases, the MSE grows linearly with s across a range of different p, which is consistent with the theoretical results.

3.3 Comparison of compressed sensing to downsampling

When designing imaging systems, one might choose between a compressed sensing setup with incoherent projections and a "downsampling" setup where we directly measure a low-resolution version of our signal

(this is specified below). The rates in Section 2.1 and 2.2 are based on assuming a variant of a sensing matrix which satisfies the RIP; hence standard CS schemes are capable of achieving the reported rates. However, in practical settings one may experience more success with a downsampling scheme. Why is this?

The key point is that many practical signals and images have significant low-resolution content or are smooth. That is, they are not only sparse, but the non-zero coefficients are very likely to correspond to low-frequency or coarse-scale information. This structure in the sparse support is not reflected by our analysis. When that structure is present, however, simple downsampling schemes can significantly outperform CS schemes, especially at low intensity levels.

We define one naïve sensing method which uses the smoothness of the signal as follows:

Definition 3.1 (Downsampling method)**.** *Let $\kappa \triangleq p/n$ be the downsampling factor (assumed to be integer-valued). The sensing matrix for downsampling is*

$$A^{\mathrm{DS}} = \mathbb{I}_{p/\kappa} \otimes \mathbb{1}_{1\times\kappa}, \tag{11}$$

where $\mathbb{I}_{p/\kappa}$ is a $p/\kappa \times p/\kappa$ identity matrix, and $\otimes$ is the Kronecker product. (That is, the first row of A^{DS} is the sum of the first κ rows of $\mathbb{I}_p$, the second row of A^{DS} is the sum of the second κ rows of $\mathbb{I}_p$, etc. The downsampling estimator for $\theta^* \triangleq D^T f^*$ is

$$\widehat{\theta}^{\mathrm{DS}} = \begin{bmatrix} \frac{1}{\sqrt{p}} \\ \\ \frac{1}{\kappa T}\bar{D}^T A^T y^{\mathrm{DS}} \end{bmatrix}. \tag{12}$$

This downsampling method enhances the SNR of any individual observation by sacrificing detailed information in the signal. Note for the general setting of compressed sensing method, knowing that the signal is smooth does little, if anything, to help the reconstruction. The compressed sensing method does not distinguish between spatial frequencies, so the reconstruction of each coefficient is equally hard. The downsampling method makes recovering low-frequency components significantly easier than recovering others, and we can obtain better performance than compressed sensing when the signal is smooth.

The theoretical bounds for the downsampling method and the details of the simulation setup can be found in Appendix B.3. In Figure 5 we compare the performance of the downsampling method and the compressed sensing method both theoretically and experimentally. In the plots, $s' \leq s$ reflects the smoothness of the signal (*i.e.,* the number of low-frequency or coarse-scale non-zero coefficients). The results show that in the low-intensity regime, when the true signal is smooth, the downsampling method can achieve better performance than the compressed sensing method. As predicted by our theory, however, once the intensity exceeds a critical threshold, we can recover both high- and low-frequency nonzero coefficients and we begin to see an improvement of the compressed sensing method over the downsampling method.

3.4 Related work

Note that this model is different from a generalized linear model of the form

$$y \sim \mathrm{Poisson}(e^{Af^*}), \tag{13}$$

where the exponentiation is considered to be element-wise. The model in (13) is appropriate for some problems and has been analyzed elsewhere in the CS and LASSO literature (*cf.,* [27, 28, 29]), but is not a good representation in the motivating applications described in the introduction. For instance, in the context of imaging, the pixel intensities in f^* are linearly combined by the physics of the imaging system, and that linear structure is not preserved by the exponentiation in (13). Furthermore, in the context of (1) and the above motivating applications, we face physical constraints on A and f^* which are not necessary in the GLM framework of (13).

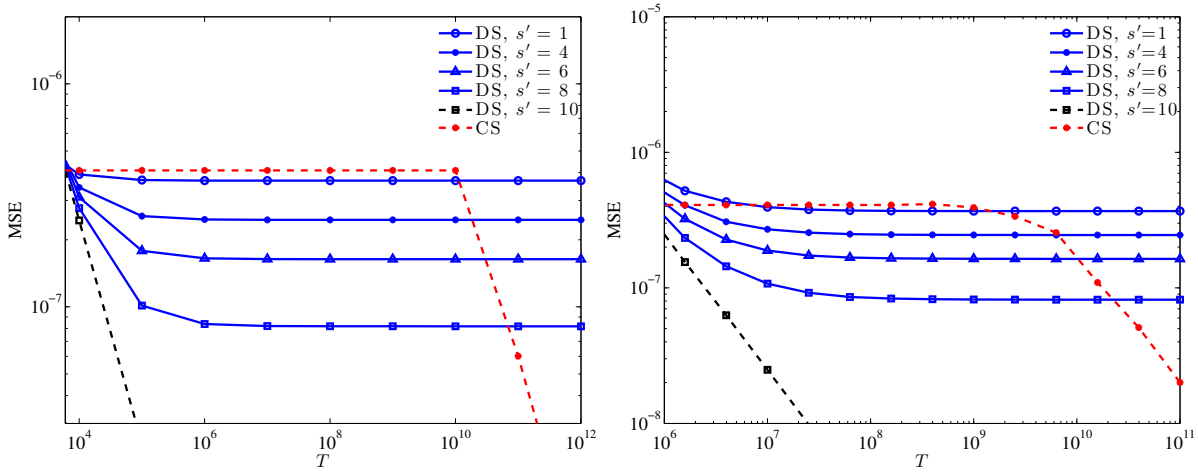

(a) Theoretical rates (b) Empirical rates

Figure 5: Theoretical and empirical rates of downsampling and compressed sensing methods. $p = 2048, K = 4, n = p/K$, and $s = 10$. s' is the number of coarse-scale nonzero coefficients which are directly measured by the proposed downsampling scheme. We see that at low-intensities, downsampling can yield much lower MSEs, but after the intensity exceeds a critical threshold, compressed sensing methods are able to estimate all nonzero coefficients accurately and the MSE is better than for downsampling schemes. This effect is predicted by our theory.

In previous performance analyses of Poisson compressed sensing, [30] provides upper bounds of reconstruction performance for a constrained ℓ_1-regularized estimator based on the restricted isometry property (RIP), [29] provides upper bounds for ℓ_1-regularized maximum likelihood estimators based on restricted eigenvalue (RE) conditions, and [31] focuses on understanding the impact of model mismatch and heteroscedasticity within standard LASSO algorithm, but none of these account the physical constraints on the sensing systems or provides lower bounds. [13, 1] consider Poisson compressed sensing with similar physical constraints to this paper, but only provided upper bounds on the risk. Furthermore, in [13], the sensing matrix is based on expander graphs, yielding different bounding techniques and algorithms. In this paper we prove both upper and lower bounds and show that they are tight in high-intensity settings.

4 Proofs

In this section we provide the proofs of Theorems 2.1 and 2.2.

4.1 Proof of Theorem 2.1

This proof of Theorem 2.1 roughly follows standard techniques for proving lower bounds on minimax rates for sparse high-dimensional problems developed in [19]. The proof involves constructing a packing set for $\mathcal{F}_{p,s,D}$ and then applying the generalized Fano method to the packing set (see e.g. [22, 23, 32] for details). The main challenge and novelty in the proof is adapting these standard arguments to the Poisson setting and constructing a packing set that satisfies all the physical constraints imposed in $\mathcal{F}_{p,s,D}$ which results in the term $\frac{k}{p^2\lambda_k^2}$.

Proof. We first introduce the packing sets that will be used in the proof. For $k = 1, \ldots, s$, let

$$\bar{\mathcal{H}}_k \triangleq \left\{\beta \in \{-1, 0, +1\}^{p-1} : \|\beta\|_0 = k\right\}.$$

By [19, Lemma 4], there exists a subset $\widetilde{\mathcal{H}}_k \subseteq \bar{\mathcal{H}}_k$ with cardinality $|\widetilde{\mathcal{H}}_k| \geq \exp\left(\frac{k}{2}\log\frac{p-k-1}{k/2}\right)$ such that the

Hamming distance $\rho_H(\beta, \beta') \geq \frac{k}{2}$ for all $\beta, \beta' \in \widetilde{\mathcal{H}}_k$. For a given $\alpha_k \in \mathbb{R}_+$, we rescale the elements of $\widetilde{\mathcal{H}}_k$ by α_k and concatenate each vector in the set with an extra entry to form the new set $\mathcal{H}_{k,\alpha_k}$:

$$\mathcal{H}_{k,\alpha_k} \triangleq \left\{\theta \in \mathbb{R}^p : \theta = [1/\sqrt{p}, \alpha_k \beta^\top]^\top, \beta \in \widetilde{\mathcal{H}}\right\}. \tag{14}$$

Define $\eta_{\alpha_k}^2 \triangleq \frac{k}{2}{\alpha_k}^2$, then the elements of $\mathcal{H}_{k,\alpha_k}$ form a η_{α_k}-packing set of $\mathcal{F}_{p,s,D}$ in the ℓ_2 norm. The following lemma describes several useful properties of this packing set.

Lemma 4.1. *For any $k = 1, \ldots, s$, let $\lambda_k = \lambda_k(\bar{D})$. Then, the packing sets $\mathcal{H}_{k,\alpha_k}$ with $0 < \alpha_k \leq \frac{1}{p\lambda_k}$ have the following properties:*

1. *The ℓ_2 distance between any two points θ and θ' in $\mathcal{H}_{k,\alpha_k}$ is bounded:*

$$\eta_{\alpha_k}^2 \leq \|\theta - \theta'\|_2^2 \leq 8\eta_{\alpha_k}^2.$$

2. *For any $\theta \in \mathcal{H}_{k,\alpha_k}$, the corresponding $f = D\theta$ satisfies:*

$$f_i \geq 0, \ \forall\, i \in \{1, \ldots, p\} \qquad \textit{and} \qquad \|f\|_1 = 1.$$

3. *The size of the packing set*

$$|\mathcal{H}_{k,\alpha_k}| \geq \exp\left(\frac{k}{2}\log\frac{p-k-1}{k/2}\right).$$

The proof of this lemma is provided in Section 4.1.1.

Throughout the proof, we also use $\Phi \triangleq AD$. Thus $\Phi\theta = Af$ when $f = D\theta$. The next several steps follow the techniques developed by [22, 23, 32]; we describe all steps for completeness. Let $M_k \triangleq |\mathcal{H}_{k,\alpha_k}|$ be the cardinality of $\mathcal{H}_{k,\alpha_k}$, and let the elements of $\mathcal{H}_{k,\alpha_k}$ be denoted $\{\theta^1, \ldots, \theta^{M_k}\}$. Define a random vector $\widetilde{\Theta} \in \mathbb{R}^p$ that is drawn from a uniform distribution over the packing set, and form a multi-way hypothesis testing problem where $\widetilde{\theta}$ is the testing result that takes value in the packing set. Then we can bound the minimax estimation error according to [32]:

$$\min_f \max_{f^* \in \mathcal{F}_{p,s,D}} \mathbb{E}\|f - f^*\|_2^2 \geq \frac{1}{4}\eta_{\alpha_k}^2 \min_{\widetilde{\theta}} \mathbb{P}[\widetilde{\theta} \neq \widetilde{\Theta}]. \tag{15}$$

By Fano's inequality and the convexity of mutual information, we have

$$\mathbb{P}[\widetilde{\theta} \neq \widetilde{\Theta}] \geq 1 - \frac{I(y; \widetilde{\Theta}) + \log 2}{\log M_k}, \tag{16}$$

where $y \sim \text{Poisson}(T\Phi\widetilde{\Theta})$ is the Poisson observation, $I(y; \widetilde{\Theta})$ is the mutual information between random variables y and $\widetilde{\Theta}$. Additionally, the mutual information can be bounded by the average of the K-L divergence between $p(y|T\Phi\theta^i)$ and $p(y|T\Phi\theta^j)$ for all $\theta^i, \theta^j \in \mathcal{H}_{k,\alpha_k}$,*i.e.,*

$$I(y; \widetilde{\Theta}) \leq \frac{1}{\binom{M_k}{2}} \sum_{\substack{i,j=1,\ldots,M_k \\ i \neq j}} \text{KL}(p(y|T\Phi\theta^i) \| p(y|T\Phi\theta^j)), \tag{17}$$

according to [22].

The following lemma provides an upper bound for the K-L divergence of Poisson distributions in terms of the squared ℓ_2-distance.

Lemma 4.2. *Let $p(y|\mu)$ denote the multivariate Poisson distribution with mean parameter $\mu \in \mathbb{R}^n_+$. For $\mu_1, \mu_2 \in \mathbb{R}^n_+$, if there exists some value $c > 0$ such that $\mu_2 \succeq c\mathbb{1}_{n\times 1}$, then the following holds:*

$$\mathrm{KL}(p(y|\mu_1)\|p(y|\mu_2)) \leq \frac{1}{c}\|\mu_1 - \mu_2\|_2^2.$$

The proof can be found in Section 4.1.2. Assumption 2.1 ensures Af^* is bounded, as we have the following lemma:

Lemma 4.3. *If the sensing matrix A satisfies Assumption 2.1, then for all nonnegative f s.t. $\|f\|_1 = 1$, we have*

$$\frac{1}{2n}\mathbb{1}_{n\times 1} \preceq Af \preceq \frac{1}{n}\mathbb{1}_{n\times 1}. \tag{18}$$

The proof is provided in Section 4.1.3.

By Lemmas 4.1 and 4.3, we have $\Phi\theta^j \succeq \frac{1}{2n}\mathbb{1}_{p\times 1}$, *i.e.,* $T\Phi\theta^j \succeq \frac{T}{2n}\mathbb{1}_{p\times 1}$. Thus by Lemma 4.2,

$$\begin{aligned} KL(p(y|T\Phi\theta^i)\|p(y|T\Phi\theta^j)) &\leq \frac{2n}{T}\|T\Phi(\theta^i - \theta^j)\|_2^2 \\ &= 2nT\|\Phi(\theta^i - \theta^j)\|_2^2. \end{aligned}$$

Let $f^i \triangleq D\theta^i$, and $f^j \triangleq D\theta^j$, then

$$\begin{aligned} \|\Phi(\theta^i - \theta^j)\|_2^2 &= \|A(f^i - f^j)\|_2^2 \\ &= \left\| \frac{1}{2(a_u - a_\ell)\sqrt{n}}\widetilde{A}(f^i - f^j) - \frac{\frac{a_u - 2a_\ell}{\sqrt{n}}}{2(a_u - a_\ell)\sqrt{n}}\mathbb{1}_{n\times p}(f^i - f^j) \right\|_2^2 \\ &= \left\| \frac{1}{2(a_u - a_\ell)\sqrt{n}}\widetilde{A}D(\theta^i - \theta^j) \right\|_2^2 \\ &\leq \frac{(1+\delta_s)}{4(a_u - a_\ell)^2 n}\|(\theta^i - \theta^j)\|_2^2. \end{aligned}$$

where the last equality is a result of $\|f^i\|_1 = \|f^j\|_1$ which leads to $\mathbb{1}_{n\times p}(f^i - f^j) = 0$, and the inequality follows from Assumption 2.2. The construction of our packing set ensures that $\|\theta^i - \theta^j\|_2^2 \leq 8\eta_{\alpha_k}^2$, so

$$\|\Phi(\theta^i - \theta^j)\|_2^2 \leq \frac{2(1+\delta_s)}{(a_u - a_\ell)^2 n}\eta_{\alpha_k}^2.$$

As a result,

$$\begin{aligned} KL(p(y|T\Phi\theta^i)\|p(y|T\Phi\theta^j)) &\leq 2nT\|\Phi(\theta^i - \theta^j)\|_2^2 \\ &\leq \frac{4(1+\delta_s)T}{(a_u - a_\ell)^2}\eta_{\alpha_k}^2. \end{aligned} \tag{19}$$

By Ineq. [(17), (19)], we can bound the mutual information

$$\begin{aligned} I(y;\widetilde{\Theta}) &\leq \frac{1}{\binom{M}{2}}\sum_{i\neq j} KL(p(y|T\Phi\theta^i)\|p(y|T\Phi\theta^j)) \\ &\leq \max_{i\neq j} KL(p(y|T\Phi\theta^i)\|p(y|T\Phi\theta^j)) \\ &\leq \frac{4(1+\delta_s)T}{(a_u - a_\ell)^2}\eta_{\alpha_k}^2. \end{aligned} \tag{20}$$

Then, by Ineq. [(16), (20)] we can bound the probability of classification error by

$$\begin{aligned}\mathbb{P}[\widetilde{\theta} \neq \widetilde{\Theta}] &\geq 1 - \frac{I(y;\widetilde{\Theta}) + \log 2}{\log M_k} \\ &\geq 1 - \frac{\frac{4(1+\delta_s)T}{(a_u - a_\ell)^2}\eta_{\alpha_k}^2 + \log 2}{\frac{k}{2}\log\frac{p-k-1}{k/2}}.\end{aligned} \tag{21}$$

We need to ensure this probability is bounded below by a positive constant; in the below we use the constant $1/4$. This constant is guaranteed as long as we have

$$\frac{k}{2}\log\frac{p-k-1}{k/2} \geq 2\log 2 \tag{22}$$

and

$$\frac{k}{2}\log\frac{p-k-1}{k/2} \geq \frac{16(1+\delta_s)T}{(a_u - a_\ell)^2}\eta_{\alpha_k}^2, \tag{23}$$

because then we can bound

$$\begin{aligned}\mathbb{P}[\widetilde{\theta} \neq \widetilde{\Theta}] &\geq 1 - \frac{\frac{4(1+\delta_s)T}{(a_u - a_\ell)^2}\eta_{\alpha_k}^2}{\frac{k}{2}\log\frac{p-k-1}{k/2}} - \frac{\log 2}{\frac{k}{2}\log\frac{p-k-1}{k/2}} \\ &\geq 1 - \frac{1}{4} - \frac{1}{2} \\ &\geq \frac{1}{4}.\end{aligned}$$

First we verify Ineq. (22). For $k \geq 2$,

$$\begin{aligned}\frac{k}{2}\log\frac{p-k-1}{k/2} &\geq \log\frac{p-k-1}{k/2} \\ &\geq \log\frac{p-s-1}{s/2} \\ &> \log\frac{p-(s+1)}{(s+1)/2} \\ &\geq \log\frac{2(s+1)}{s/2} \\ &= 2\log 2,\end{aligned}$$

where the last inequality is a result of $p \geq 3(s+1)$. For $k = 1$,

$$\begin{aligned}\frac{k}{2}\log\frac{p-k-1}{k/2} &= \frac{1}{2}\log\frac{p-2}{1/2} \\ &\geq \frac{1}{2}\log\frac{10-2}{1/2} \\ &= 2\log 2,\end{aligned}$$

where the inequality is a result of $p \geq 10$. Now to ensure Ineq. (23), we need

$$\frac{(1+\delta_s)T}{(a_u - a_\ell)^2}\eta_{\alpha_k}^2 \leq \frac{k}{32}\log\frac{p-k-1}{k/2},$$

which leads to

$$\eta_{\alpha_k}^2 \leq \frac{(a_u - a_\ell)^2 k}{32(1+\delta_s)T} \log \frac{p-k-1}{k/2}.$$

Next recall the condition of Lemma 4.1 that $0 < \alpha_k \leq \frac{1}{p\lambda_k}$. This combined with the above suggests α_k must be chosen so that

$$\eta_{\alpha_k}^2 = \min \left\{ \frac{k}{2} \left(\frac{1}{p\lambda_k} \right)^2, \frac{(a_u - a_\ell)^2 k}{32(1+\delta_s)T} \log \frac{p-k-1}{k/2} \right\},$$

which is satisfied by our choice of α_k. A lower bound is then proved using the packing set $\mathcal{H}_{\alpha_k}$ such that

$$\begin{aligned} \min_f \max_{f^* \in \mathcal{F}_{p,s,D}} \mathbb{E}\|f - f^*\|_2^2 &\geq \frac{1}{4}\eta_{\alpha_k}^2 \mathbb{P}[\widetilde{\theta} \neq \widetilde{\Theta}] \geq \frac{1}{16}\eta_{\alpha_k}^2 \\ &\geq \min \left\{ \frac{1}{32} \frac{k}{p^2\lambda_k^2}, \frac{(a_u - a_\ell)^2 k}{512(1+\delta_s)T} \log \frac{p-k-1}{k/2} \right\}. \end{aligned}$$

For each level $k = 1, \ldots, s$, the above is a separate lower bound for the minimax risk. To make the bound as tight as possible, we take the maximum of the bounds for all s sparsity levels, which results in

$$\min_f \max_{f^* \in \mathcal{F}_{p,s,D}} \mathbb{E}\|f - f^*\|_2^2 \geq \max_{1 \leq k \leq s} \left\{ \min \left(\frac{1}{32} \frac{k}{p^2\lambda_k^2}, \frac{(a_u - a_\ell)^2 k}{512(1+\delta_s)T} \log \frac{p-k-1}{k/2} \right) \right\}.$$

Then there exists some constant $C_L > 0$ such that

$$\min_f \max_{f^* \in \mathcal{F}_{p,s,D}} \mathbb{E}\|f - f^*\|_2^2 \geq C_L \max_{1 \leq k \leq s} \left\{ \min \left(\frac{k}{p^2\lambda_k^2}, \frac{k}{(1+\delta_s)T} \log \frac{p-k-1}{k/2} \right) \right\},$$

which completes the proof. □

4.1.1 Proof of Lemma 4.1

Proof. We first prove property 1. Let $\bar{\theta} = [\theta_2, \ldots, \theta_p]^\top, \bar{\theta}' = [\theta_2', \ldots, \theta_p']^\top$. Note that the first elements of θ and θ' are equal, *i.e.,* $\theta_1 = \theta_1'$, thus $\|\theta - \theta'\|_2^2 = \|\bar{\theta} - \bar{\theta}'\|_2^2$. By construction, the Hamming distance between θ and θ' is at least $k/2$. Because $\bar{\theta}, \bar{\theta}' \in \{0, \pm\alpha_k\}^{p-1}$, we have

$$\|\bar{\theta} - \bar{\theta}'\|_2^2 \geq \frac{k}{2}\alpha_k^2 = \eta_{\alpha_k}^2.$$

Next, note that $\|\bar{\theta}\|_0 = \|\bar{\theta}'\|_0 = k$. Thus the maximum distance is

$$\|\bar{\theta} - \bar{\theta}'\|_2^2 \leq k(\alpha_k - (-\alpha_k))^2 = 4k\alpha_k^2 = 8\eta_{\alpha_k}^2,$$

and we have property 1 proved.

For property 2, we have that for any $\theta \in \mathcal{H}_{k,\alpha_k}$, $\theta_1 = 1/\sqrt{p}$, and $|\theta_i| \leq \alpha_k \leq \frac{1}{p\lambda_k}, i \neq 1$. Let $\bar{\theta} = [\theta_2, \ldots, \theta_p]^\top$ and $\beta = \bar{\theta}/\alpha_k \in \{0, \pm 1\}^{p-1}$ be the corresponding element in $\bar{\mathcal{H}}_k$, then

$$\begin{aligned} f_i &= (D\theta)_i \\ &= D_{i,1}\theta_1 + \sum_{j=2}^{p} D_{i,j}\theta_j \\ &= \frac{1}{\sqrt{p}}\theta_1 + \alpha_k \langle \bar{D}_i, \beta \rangle \\ &\geq \frac{1}{\sqrt{p}}\frac{1}{\sqrt{p}} - \alpha_k |\langle \bar{D}_i, \beta \rangle| \\ &\geq \frac{1}{p} - \alpha_k \lambda_k. \end{aligned}$$

By $\alpha_k \leq \frac{1}{p\lambda_k}$,

$$\begin{aligned} f_i &\geq \frac{1}{p} - \alpha_k \lambda_k \\ &\geq \frac{1}{p} - \frac{1}{p\lambda_k}\lambda_k \\ &= 0. \end{aligned}$$

Then, by the properties of D:

$$\|D\theta\|_1 = \sqrt{p}\theta_1 = 1.$$

Property 3 is satisfied by construction. The proof of the size of the packing set can be found in [19, Lemma 4]. □

4.1.2 Proof of Lemma 4.2

Proof. The proof is straight-forward. For Poisson distributions, we have

$$\begin{aligned} KL(p(y|\mu_1)\|p(y|\mu_2)) &= \sum_{i=1}^{n}\left[(\mu_1)_i \log\frac{(\mu_1)_i}{(\mu_2)_i} - (\mu_1)_i + (\mu_2)_i\right] \\ &\leq \sum_{i=1}^{n}\left[(\mu_1)_i\left(\frac{(\mu_1)_i}{(\mu_2)_i} - 1\right) - (\mu_1)_i + (\mu_2)_i\right] \\ &= \sum_{i=1}^{n}\frac{1}{(\mu_2)_i}\left[(\mu_1)_i^2 - 2(\mu_1)_i(\mu_2)_i + (\mu_2)_i^2\right] \\ &\leq \frac{1}{\min_i(\mu_2)_i}\sum_{i=1}^{n}\left[(\mu_1)_i^2 - 2(\mu_1)_i(\mu_2)_i + (\mu_2)_i^2\right] \\ &\leq \frac{1}{c}\|\mu_1 - \mu_2\|_2^2, \end{aligned}$$

where the last inequality is a result of $\mu_2 \succeq c\mathbb{1}_{n\times 1}$. □

4.1.3 Proof of Lemma 4.3

Proof. For any A satisfying Assumption 2.1, by Eq. 40, we have $A_{i,j} \in [\frac{1}{2n}, \frac{1}{n}]$. Thus for any nonnegative f satisfying $\|f\|_1 = 1$, by Hölder's inequality for $i = 1, \ldots, p$, we have

$$(Af)_i = \sum_{j=1}^{p} A_{i,j}f_j \geq \|f\|_1 \min_{j=1,\ldots,p} A_{i,j} \geq \frac{1}{2n},$$

and

$$(Af)_i = \sum_{j=1}^{p} A_{i,j}f_j \leq \|f\|_1 \max_{j=1,\ldots,p} A_{i,j} \leq \frac{1}{n}.$$

□

4.2 Proof of Theorem 2.2

Proof. Here we use the proof technique developed in [1], but modify it to provide sharper upper bounds. First, we do not assume f^* bigger than some positive threshold $c/p\mathbb{1}_{p\times 1}$. Instead, we work with a more general set of signals where we only assume f^* is nonnegative. We also modify the feasible set of estimators and use slightly different assumptions on the sensing matrix A to further tighten the bound. The main steps of the proof include:

1. Construct a specific finite set Γ of quantized feasible estimators, and design a penalty function $\mathrm{pen}(\cdot)$;
2. Use the assumptions on A and the Li-Barron theorem on the error of penalized likelihood estimators (see [1, 33, 34]) to bound

$$\|\widehat{f} - f^*\|_2^2 \le c \min_{(D^\top f)\in\Gamma} \left\{ \|f - f^*\|_2^2 + \frac{1}{T}\,\mathrm{pen}(f) \right\},$$

   for some constant c;
3. Choose one estimator in the feasible set, and compute its cost as the upper bound.

Let $L \triangleq \max_{j,k} |D_{j,k}|$ and note

$$\|\theta^*\|_\infty = \|D^\top f^*\|_\infty \le \max_{j,k} |D_{j,k}| \|f^*\|_1 = \max_{i,j} |D_{i,j}| = L.$$

Next define the quantization level

$$\begin{aligned} K^* &\triangleq \sqrt{\frac{(1+\delta_s)L^2 T \log 2}{2(a_u - a_\ell)^2}}, \\ K &\triangleq \lceil K^* \rceil . \end{aligned} \tag{24}$$

Let

$$\mathcal{B} \triangleq \begin{cases} \{0\}, & K = 1 \\ \left\{-L, -\frac{K-3}{K-1}L, -\frac{K-5}{K-1}L, \ldots, \frac{K-3}{K-1}L, L\right\}, & K > 1, \end{cases}$$

where $|\mathcal{B}| = K$. We now consider the case where $K > 1$ (which also implies $K^* > 1$); the case where $K = 1$ will be addressed later. Define the following sets:

$$\Theta_K \triangleq \left\{ \theta \in \mathbb{R}^p : \theta_1 = p^{-1/2}, \|\theta\|_0 \le s+1, \theta_i \in \mathcal{B} \text{ for } i = 2, \ldots, p \right\}, \tag{25a}$$

$$\mathcal{F} \triangleq \{ f = D\theta : \theta \in \Theta_K \}, \tag{25b}$$

$$\mathcal{C} \triangleq \left\{ g \in \mathbb{R}_+^p : \|g\|_1 = 1 \right\}, \tag{25c}$$

$$\Gamma \triangleq \left\{ \widetilde{\theta} = D\widetilde{f} : f \in \mathcal{F}, \widetilde{f} = \operatorname*{argmin}_{g\in\mathcal{C}} \|f - g\|_2 \right\}. \tag{25d}$$

Further define the penalty function

$$\mathrm{pen}(\theta) = 2\log_2(s) + 2\|\bar{\theta}\|_0 \log_2(pK). \tag{26}$$

This penalty function is calculated from the code length of a prefix code for all $\theta \in \Theta_K$. Since we know θ_1 exactly, we only need to code $\bar{\theta} = [\theta_2, \ldots, \theta_p]^\top$. The prefix code can be generated using the following steps:

1. First, we encode $\|\bar{\theta}\|_0$, the number of non-zeros of $\bar{\theta}$, which requires $\log_2 s$ bits.
2. Second, we encode each of the locations of the $\|\bar{\theta}\|_0$ non-zero elements. Each location takes $\log_2 p$ bits, so the overall length is $\|\bar{\theta}\|_0 \log_2 p$ bits.
3. Last, we encode the value of each non-zero element. Since each value is quantized to one of the K bins, this step will take $\|\bar{\theta}\|_0 \log_2 K$ bits.

Thus the prefix code length is $\log_2(s) + \|\bar{\theta}\|_0 \log_2(pK)$. Note our penalty function is twice the code length. This scaling is needed later for the proof of Eq. (30). By having a corresponding prefix code, the penalty function (rescaled by $1/2$) satisfies the Kraft inequality: $\sum_{\theta \in \Theta_K} e^{-\operatorname{pen}(\theta)/2} \leq 1$. This code can also be applied to the set Γ by coding $\widetilde{\theta}$ using the same code for the corresponding $\theta \in \Theta_K$, where we also have

$$\sum_{\widetilde{\theta} \in \Gamma} e^{-\operatorname{pen}(\widetilde{\theta})/2} \leq 1. \tag{27}$$

By Assumption 2.3 and $\|f^*\|_1 = \|\widehat{f}\|_1$, we have

$$\begin{aligned} \|f^* - \widehat{f}\|_2^2 &\leq \frac{1}{(1-\delta_s)} \|\widetilde{A}(f^* - \widehat{f})\|_2^2 \\ &\leq \frac{(2(a_u - a_\ell)\sqrt{n})^2}{(1-\delta_s)} \|A(f^* - \widehat{f})\|_2^2. \end{aligned} \tag{28}$$

Then,

$$\begin{aligned} \|A(f^* - \widehat{f})\|_2^2 &= \sum_{i=1}^n \left[(Af^*)_i - (A\widehat{f})_i \right]^2 \\ &= \sum_{i=1}^n \left[(Af^*)_i^{1/2} - (A\widehat{f})_i^{1/2} \right]^2 \left[(Af^*)_i^{1/2} + (A\widehat{f})_i^{1/2} \right]^2 \\ &\leq 4 \max_i \{ (Af^*)_i, (A\widehat{f})_i \} \sum_{i=1}^n \left[(Af^*)_i^{1/2} - (A\widehat{f})_i^{1/2} \right]^2 \\ &\leq \frac{4}{n} \sum_{i=1}^n \left[(Af^*)_i^{1/2} - (A\widehat{f})_i^{1/2} \right]^2, \end{aligned} \tag{29}$$

where the second equality holds because of the non-negativity constraints. The last inequality holds because $D^\top \widehat{f} \in \Gamma \Rightarrow \widehat{f} \in \mathcal{C} \Rightarrow \widehat{f}$ is nonnegate and $\|\widehat{f}\|_1 = 1$, so we can apply Lemma 4.3 to both $\widehat{f}$ and f^*.

Following the technique used in [1, Eq. (25) and Ineq. (26)] (where we need the Kraft inequality in (27)), we further have

$$\mathbb{E}\left[\sum_{i=1}^n \left((TAf^*)_i^{1/2} - (TA\widehat{f})_i^{1/2} \right)^2 \right] \leq \min_{(D^\top f) \in \Gamma} \left[\mathrm{KL}\left(p(\cdot|TAf^*) \| p(\cdot|TAf) \right) + \operatorname{pen}(f) \right], \tag{30}$$

which is equivalent to

$$\mathbb{E}\left[\sum_{i=1}^n \left((Af^*)_i^{1/2} - (A\widehat{f})_i^{1/2} \right)^2 \right] \leq \frac{1}{T} \min_{(D^\top f) \in \Gamma} \left[\mathrm{KL}\left(p(\cdot|TAf^*) \| p(\cdot|TAf) \right) + \operatorname{pen}(f) \right].$$

By Lemma 4.2 and $Af \succeq \frac{1}{2n} \mathbb{1}_{n \times 1}$ from Lemma 4.3, we have

$$\mathrm{KL}\left(p(\cdot|TAf^*) \| p(\cdot|TAf) \right) \leq \frac{2n}{T} \|TA(f^* - f)\|_2^2 = 2nT \|A(f^* - f)\|_2^2.$$

Thus,

$$\mathbb{E}\left[\sum_{i=1}^n \left((Af^*)_i^{1/2} - (A\widehat{f})_i^{1/2} \right)^2 \right] \leq \min_{(D^\top f) \in \Gamma} \left[2n \|A(f^* - f)\|_2^2 + \frac{1}{T} \operatorname{pen}(f) \right]. \tag{31}$$

Now we can apply Assumption 2.2 and have

$$\|A(f^* - f)\|_2^2 = \frac{1}{(2(a_u - a_\ell)\sqrt{n})^2}\|\widetilde{A}(f^* - f)\|_2^2 \le \frac{(1+\delta_s)}{4(a_u - a_\ell)^2 n}\|f^* - f\|_2^2, \quad \forall (D^\top f) \in \Gamma. \tag{32}$$

Combining Eq. (28), (29), (31), and (32), we get

$$\begin{aligned}\mathbb{E}\left[\|f^* - \widehat{f}\|_2^2\right] &\le \frac{(2(a_u - a_\ell)\sqrt{n})^2}{(1-\delta_s)} \cdot \frac{4}{n} \min_{(D^\top f)\in\Gamma}\left[\frac{(1+\delta_s)}{2(a_u - a_\ell)^2}\|f^* - f\|_2^2 + \frac{1}{T}\operatorname{pen}(f)\right] \\ &= \min_{(D^\top f)\in\Gamma}\left[\frac{8(1+\delta_s)}{(1-\delta_s)}\|f^* - f\|_2^2 + \frac{16(a_u - a_\ell)^2}{(1-\delta_s)T}\operatorname{pen}(f)\right].\end{aligned} \tag{33}$$

We now bound $\|f^* - f\|_2^2$ and $\operatorname{pen}(f)$ in (33) for some good choice of f. In particular, we choose f_q such that $D^T f_q \in \Theta_K$ is the quantized version of $\theta^* = D^\top f^*$; *i.e.,* θ_q is the projection of θ^* onto Θ_K. For any $f_q \in \mathcal{F}$, let $\widetilde{f}_q \in \mathcal{C}$ be the projection of f_q onto $\mathcal{C}$, and we have $D^\top \widetilde{f}_q \in \Gamma$. By the Pythagorean identity we have $\|f^* - \widetilde{f}_q\|_2^2 \le \|f^* - f_q\|_2^2$ because $f^* \in \mathcal{C}$ and $\widetilde{f}_q$ is the projection of f_q onto $\mathcal{C}$. Furthermore, we have $\operatorname{pen}(\widetilde{f}_q) = \operatorname{pen}(f_q)$ because we use the same code for $f_q \in \mathcal{F}$ and its projection $\widetilde{f}_q \in \mathcal{C}$. Thus

$$\mathbb{E}\left[\|f^* - \widehat{f}\|_2^2\right] \le \frac{8(1+\delta_s)}{(1-\delta_s)}\|f^* - f_q\|_2^2 + \frac{16(a_u - a_\ell)^2}{(1-\delta_s)T}\operatorname{pen}(f_q). \tag{34}$$

We now bound $\|f^* - f_q\|_2^2$ in (34). Since $\theta_q \in \Theta_K$ is the $(s+1)$-sparse quantized version of θ^* so that $(\theta_q)_1 = \theta_1^*$ and $|(\theta_q)_i - \theta_i^*| \le \frac{L}{K}, i = 2, \ldots, p$, we have

$$\|f^* - f_q\|_2^2 = \|\theta^* - \theta_q\|_2^2 \le s\left(\frac{L}{K}\right)^2 = \frac{sL^2}{K^2}. \tag{35}$$

Using (26), we have

$$\operatorname{pen}(f_q) = 2\log_2 s + 2s\log_2 p + 2s\log_2 K.$$

Thus we can bound

$$\mathbb{E}\left[\|f^* - \widehat{f}\|_2^2\right] \le \frac{8(1+\delta_s)}{(1-\delta_s)}\frac{sL^2}{K^2} + \frac{32(a_u - a_\ell)^2}{(1-\delta_s)T}[\log_2 s + s\log_2 p + s\log_2 K], \tag{36}$$

which depends on K. The value of K^* defined in (24) minimizes this bound, and note that $1 < K* \le K \le K^* + 1$. Thus

$$\begin{aligned}\mathbb{E}\left[\|f^* - \widehat{f}\|_2^2\right] &\le \frac{8(1+\delta_s)}{(1-\delta_s)}\frac{sL^2}{(K^*)^2} + \frac{32(a_u - a_\ell)^2}{(1-\delta_s)T}[\log_2 s + s\log_2 p + s\log_2(K^* + 1)] \\ &\le \frac{8(1+\delta_s)}{(1-\delta_s)}\frac{sL^2}{(K^*)^2} + \frac{32(a_u - a_\ell)^2}{(1-\delta_s)T}[\log_2 s + s\log_2 p + s\log_2\left((K^*)^2 + 1\right)].\end{aligned}$$

Note that $L \le 1$ because D is orthonormal, and using the value for K^* and L yields

$$\begin{aligned}\mathbb{E}\left[\|f^* - \widehat{f}\|_2^2\right] \le{}& \frac{16(a_u - a_\ell)^2 s}{(1-\delta_s)T\log 2} + \frac{32(a_u - a_\ell)^2}{(1-\delta_s)}\frac{1}{T}[\log_2 s + s\log_2 p] \\ &+ \frac{32(a_u - a_\ell)^2}{(1-\delta_s)}\frac{s}{T}\left[\log_2\left(\frac{(1+\delta_s)\log 2}{2(a_u - a_\ell)^2}T + 1\right)\right].\end{aligned}$$

Note there exists some constant $c_0 > 0$ s.t. $\log_2\left(\frac{(1+\delta_s)\log 2}{2(a_u-a_\ell)^2}T+1\right) \le c_0\log(T+1)$, thus the dominating terms in the bound are proportional to $\frac{s}{T}\log_2 p$ and $\frac{s}{T}\log_2(T+1)$. Then there exist some constants $c_1, c_2 > 0$ such that

$$\mathbb{E}[\|f^* - \widehat{f}\|_2^2] \le \frac{c_1 s\log_2 p}{(1-\delta_s)T} + \frac{c_2 s\log_2(T+1)}{(1-\delta_s)T}. \tag{37}$$

Instead of using the quantized estimator $\widehat{f}$, an alternative estimator is the mean estimator $\widehat{f} = \frac{1}{\sqrt{p}}\mathbb{1}_{n\times 1}$. This estimator corresponds to the case where $K = 1$ above.

$$\mathbb{E}[\|f^* - \widehat{f}\|_2^2] = \|f^* - \frac{1}{\sqrt{p}}\mathbb{1}_{n\times 1}\|_2^2 = \|\bar{\theta}^*\|_2^2. \tag{38}$$

$\|\bar{\theta}^*\|_2^2$ can be upper bounded in one of two ways. Either $\|\bar{\theta}^*\|_2^2 < \|f^*\|_2^2 \le \|f^*\|_1^2 = 1$, or

$$\|\bar{\theta}^*\|_2^2 \le s\|\theta^*\|_\infty^2 = s\|D^T f^*\|_\infty^2 \le s\max_{j,k}|D_{j,k}|^2\|f^*\|_1^2 = s\max_{j,k}|D_{j,k}|^2.$$

Since the minimax mean-squared error takes the minimum over all measurable estimators, we have for some $C_U > 0$,

$$\min_{\widehat{f}}\max_{f^*\in\mathcal{F}_{p,s,D}}\mathbb{E}[\|f^* - \widehat{f}\|_2^2] \quad \le C_U\max\left(1, s\max_{j,k}|D_{j,k}|^2, \tfrac{s\log_2 p}{(1-\delta_s)T} + \tfrac{s\log_2(T+1)}{(1-\delta_s)T}\right). \tag{39}$$

□

As noted earlier, the lower bound in Theorem 2.1 only requires Assumptions 2.1 and 2.2, whereas the upper bound requires the additional Assumption 2.3. To explain this, first note that by Lemma 4.3, the intensity of Af^* is both lower and upper bounded. Then with the upper bound of $\|\widetilde{A}f\|_2^2$, the KL-divergence is upper bounded by the squared ℓ_2 distance; with the lower bound of $\|\widetilde{A}f\|_2^2$, the KL-divergence is lower bounded by the squared ℓ_2 distance. The lower bound only requires that the KL-divergence between parameters in the space is upper bounded by the squared ℓ_2 distance whereas the upper bound requires that the KL-divergence be both upper and lower bounded in terms of ℓ_2 distance.

5 Conclusion

In this paper, we provide sharp reconstruction performance bounds for photon-limited imaging systems with real-world constraints. These critical physical constraints, referred to as the positivity and flux-preserving constraints, arise naturally from the fact that we cannot subtract light in an optical system, and that the total number of available photons (which is proportional to the total observation time or total source intensity) is fixed regardless of the number of sensors in a system. These constraints, which are often not considered in other literature, play an essential role in our performance guarantees.

Unlike most compressed sensing results, our mean-squared error bounds do not decrease as n, the number of observations, grows to infinity. This independence of n derives from the flux-preserving constraint. In most compressed sensing results, the signal-to-noise ratio is controlled by the number of observations, while here, as we work in a photon-limited environment, we cannot afford an infinite photon-count or observation time. In this photon-limited framework, the signal-to-noise ratio is controlled by the total intensity, T. This allows us to separate the effect of having a higher signal-to-noise ratio (larger T) and the effect of having a better variety of observations (larger n). As a result, we see that we cannot drive the reconstruction error to 0 simply by having more measurements. Instead, we need infinite signal-to-noise ratio ($T \to \infty$) to achieve perfect reconstruction. However, it is still critical to have n sufficiently large in order to have all the assumptions on the sensing matrix satisfied.

Another interesting observation is that the lower and upper bounds derived in this paper both exhibit different behaviors in the low- and high-intensity regimes. Such behavior change is also demonstrated in

the simulations, where we see "elbows" in the plot of MSE vs. T. The low-intensity region, in particular, corresponds to a range of T where reliable reconstruction is hard to achieve.

In conventional CS settings, CS does not lead to significantly higher MSEs than directly sensing nonzero coefficients (if their locations were known). In contrast, because of the unusual role of noise in our setting, directly sensing nonzero coefficients can lead to dramatic reductions in MSE. This is the reason that optical systems which measure low-resolution images directly rather than collect compressive measurements can perform so much better in practice, particularly in low-intensity regimes, as detailed in Section 3.3. For example, most night vision cameras have a focal plane array with n elements, each of which directly measures a distinct region of the field of view. Conventional CS suggests we could change the optical design so that the same n-element focal plane array could be used to generate an estimate of the scene with $p > n$ pixels. However, our results suggest that this approach will not be successful unless T (the observation/stare time or scene brightness) is quite large.

Finally, we note that our minimax upper bounds do not correspond to a feasible estimator. Developing upper bounds for estimators which are implementable in polynomial time (*e.g.,* corresponding to ℓ_1 regularization as in the LASSO) is an important avenue for future work. Because of the close correspondence between our theoretical upper bounds and our simulations (which used an ℓ_1 regularization algorithm), we anticipate that the bounds on such a method would not differ significantly from the bounds presented in this paper.

6 Acknowledgements

We gratefully acknowledge the support of the awards AFOSR FA9550-11-1-0028 and NSF CCF-06-43947. GR was partially supported from August 2012 to July 2013 by the National Science Foundation under Grant DMS-1127914 to the Statistical and Applied Mathematical Sciences Institute. We also thank Dr. Zachary Harmany for the fruitful discussions.

A Proofs

A.1 Proof of Lemma 2.1

Proof. By $\widetilde{A}_{i,j} \in [a_\ell, a_u]$, we have

$$A_{i,j} = \frac{\widetilde{A}_{i,j} + \frac{a_u - 2a_\ell}{\sqrt{n}} \mathbb{1}_{n\times p}}{2(a_u - a_\ell)\sqrt{n}} \in \left[\frac{1}{2n}, \frac{1}{n}\right]. \tag{40}$$

Thus we have,

$$\begin{aligned} A_{i,j} &\geq 0, && \forall i, j, \\ \sum_{i=1}^{n} A_{i,j} &\leq \sum_{i=1}^{n} \frac{1}{n} = 1, && \forall j. \end{aligned}$$

□

The value of shift and rescaling parameters in Eq. (3) are chosen to ensure $A_{i,j} \in [\frac{1}{2n}, \frac{1}{n}]$. The upper bound of $1/n$ is needed to satisfy the flux-preserving constraint, while the lower bound of $1/(2n)$ is chosen to ensure that that we have a reasonable lower bound on Af^* so that we can lower bound the KL divergence between y and TAf^*. It is possible to change the shift and rescaling parameters to have an arbitrary lower bound $\epsilon < 1/n$ on $A_{i,j}$. However, ϵ being either too small (≈ 0) or too big ($\approx 1/n$) will result in suboptimal bounds. The best choice of ϵ is c/n for some $0 < c < 1$, where we use $c = 1/2$ in Assumption 2.1.

A.2 Proof of Lemma 3.1

The proof technique is similar to the proof in Appendix 4.1. We omit some repeated steps for conciseness.

Proof. First define the sets

$$\mathcal{G} = \{f \in \{0,1\}^p : \|f\|_0 = 1\},$$
$$\mathcal{H} = \{\theta \in \mathbb{R}^p : \theta = D^\top f, f \in \mathcal{G}\}.$$

Note $\mathcal{G} \subset \mathcal{F}_{p,s,D}$, and for this Haar Wavelet basis, all $\theta \in \mathcal{H}$ satisfies $\|\theta\|_0 = 1 + \log_2 p$. Thus θ is $(s+1)$-sparse for all $s \geq \log_2 p$. Also we have

$$\|\theta^i - \theta^j\|_2^2 = \|f^i - f^j\|_2^2 = 2, \forall i, j = 1, \ldots, M.$$

Thus $\mathcal{H}$ is a $\sqrt{2}$-packing set of $\mathcal{F}_{p,s,D}$ in the ℓ_2 norm. The cardinality of $\mathcal{H}$ is $M \triangleq |\mathcal{H}| = |\mathcal{G}| = p$. Let the elements of $\mathcal{H}$ be denoted $\{\theta^1, \ldots, \theta^M\}$ and $f^i \triangleq D\theta^i, i = 1, \ldots, M$. Then we can define a random vector $\widetilde{\Theta} \in \mathbb{R}^p$ that is drawn from a uniform distribution over the set $\mathcal{H}$, and form a multi-way hypothesis testing problem where $\widetilde{\theta}$ is the testing result that takes value in $\mathcal{H}$. Then we can bound the minimax estimation error using the same proof technique in Appendix 4.1 and get

$$\begin{aligned}\min_f \max_{f^* \in \mathcal{F}_{p,s,D}} \mathbb{E}\|f - f^*\|_2^2 &\geq \frac{1}{2} \min_{\widetilde{\theta}} \mathbb{P}[\widetilde{\theta} \neq \widetilde{\Theta}] \\ &\geq \frac{1}{2}\left(1 - \frac{I(y; \widetilde{\Theta}) + \log 2}{\log M}\right),\end{aligned}$$

where

$$\begin{aligned}I(y; \widetilde{\Theta}) &\leq \frac{1}{\binom{M}{2}} \sum_{\substack{i,j=1,\ldots,M \\ i \neq j}} \mathrm{KL}(p(y|T\Phi\theta^i) \| p(y|T\Phi\theta^j)) \\ &\leq \frac{(1+\delta_s)T}{2(a_u - a_\ell)^2} \|(\theta^i - \theta^j)\|_2^2 \\ &= \frac{(1+\delta_s)T}{(a_u - a_\ell)^2}.\end{aligned}$$

By $\log M = \log p \geq 2 \log 2$ and $T \leq \frac{(a_u - a_\ell)^2 \log p}{4(1+\delta_s)}$, we have

$$1 - \frac{I(y; \widetilde{\Theta}) + \log 2}{\log M} \geq \frac{1}{4}.$$

Thus there exists an absolute constant $C_L'' \geq 1/8$ s.t.

$$\min_f \max_{f^* \in \mathcal{F}_{p,s,D}} \mathbb{E}\|f - f^*\|_2^2 \geq C_L''.$$

□

A.3 Proof of Lemma B.1

Proof. Recall that we know the first coefficient of θ^*. The total mean square error for the downsampling method is then

$$\begin{aligned}R^{\mathrm{DS}} &= \mathbb{E}\left[\|\widehat{\theta}^{\mathrm{DS}} - \theta^*\|_2^2\right] \\ &= \sum_{i=2}^{p/\kappa} \mathbb{E}\left[(\widehat{\theta}_i^{\mathrm{DS}} - \theta_i^*)^2\right] + \sum_{j=p/\kappa+1}^{p} \mathbb{E}\left[(\widehat{\theta}_j^{\mathrm{DS}} - \theta_j^*)^2\right].\end{aligned} \tag{41}$$

In Eq (41), $i = 2, \ldots, p/\kappa$ correspond to the coefficients in the coarse scales. For DWT, the downsampling process does not distort these coefficients, and the reconstruction is unbiased, *i.e.,*

$$\mathbb{E}\left[\widehat{\theta}_i^{\mathrm{DS}} - \theta_i^*\right] = 0, \quad i = 2, \ldots, p/\kappa.$$

Thus by bias-variance decomposition, $\sum_{i=2}^{p/\kappa} \mathbb{E}\left[(\widehat{\theta}_i^{\mathrm{DS}} - \theta_i^*)^2\right]$ only depends on the variance of the estimator $\widehat{\theta}_i^{\mathrm{DS}}$. To compute $\mathrm{Var}\left[\widehat{\theta}_i^{\mathrm{DS}}\right]$, we first evaluate $\mathrm{Var}\left[\widehat{f}_k^{\mathrm{DS}}\right]$ for $k = 1, \ldots, p$, and then use the basic property of the variance to compute the variance of $\widehat{\theta}_i^{\mathrm{DS}} = D^\top \widehat{f}^{\mathrm{DS}}$. Note that $A^\top$ has only one non-zero entry (with value 1) in each row, thus each of the estimator $\widehat{f}_k^{\mathrm{DS}}$ is a scaled Poisson random variable, whose variance is the scaled mean parameter of the Poisson process:

$$\mathrm{Var}\left[\widehat{f}_k^{\mathrm{DS}}\right] = \frac{1}{\kappa^2 T^2}(TAf^*)_{k'},$$

where $k' = \lfloor k/\kappa \rfloor + 1$, and $(Af^*)_{k'} = \sum_{j=1}^{\kappa} f^*_{(k'-1)\kappa+j}$. Note that $f_j^* \le 2/p$ because

$$\begin{aligned}
f_i^* &= (D\theta^*)_i \\
&= D_{i,1}\theta_1 + \sum_{j=2}^{p} D_{i,j}\theta_j \\
&= \frac{1}{\sqrt{p}}\theta_1 + \alpha\langle \bar{D}_i, \bar{\theta}\rangle \\
&\le \frac{1}{\sqrt{p}}\frac{1}{\sqrt{p}} + \alpha_k|\langle \bar{D}_i, \bar{\theta}\rangle| \\
&\le \frac{1}{p} + \alpha\lambda \\
&= \frac{1}{p} + \frac{1}{p\lambda}\lambda \\
&= \frac{2}{p},
\end{aligned}$$

and we get

$$\mathrm{Var}\left[\widehat{f}_k^{\mathrm{DS}}\right] \le \frac{1}{\kappa^2 T^2}\kappa\frac{2T}{p} = \frac{2}{Tp\kappa}.$$

Now by $\widehat{\theta}_i^{\mathrm{DS}} = D^\top \widehat{f}^{\mathrm{DS}}$, and D orthonormal, we have

$$\begin{aligned}
\mathrm{Var}\left[\widehat{\theta}_i^{\mathrm{DS}}\right] &= \sum_{k=1}^{p} D_{ki}^2 \mathrm{Var}\left[\widehat{f}_k^{\mathrm{DS}}\right] \\
&\le \frac{2}{Tp\kappa}\sum_{k=1}^{p} D_{ki}^2 \\
&= \frac{2}{Tp\kappa}.
\end{aligned}$$

Note we have s' coefficients θ_i that are in the coarse scale, we can conclude by

$$\sum_{i=2}^{p/\kappa} \mathbb{E}\left[(\widehat{\theta}_i^{\mathrm{DS}} - \theta_i^*)^2\right] \le \frac{2s'}{Tp\kappa}.$$

On the other hand, $j = p/\kappa + 1, \ldots, p$ correspond to the coefficients in the fine scales. For the DWT basis, the downsampling process completely erases these information because all details in the fine scales become unidentifiable after the downsampling process. As a result, $\widehat{\theta}_j^{\mathrm{DS}} = 0$, and $\mathbb{E}\left[(\widehat{\theta}_j^{\mathrm{DS}} - \theta_j^*)^2\right] = (\theta_j^*)^2 = \alpha_s^2 = \frac{1}{p^2\lambda^2}$.

Add the coarse scale and fine scale coefficients together, we can get

$$\begin{aligned} R^{\mathrm{DS}} &= \mathbb{E}\left[\|\widehat{\theta}^{\mathrm{DS}} - \theta^*\|_2^2\right] \\ &= \frac{2s'}{Tp\kappa} + \frac{s - s'}{\lambda^2 p^2}, \end{aligned}$$

which completes the proof. □

B Others

B.1 Generation of signals with different sparse supports

θ_1^* is the triangular signal, where

$$(\theta_1^*)_j = \begin{cases} \frac{1}{\sqrt{p}}, & j = 1 \\ \frac{(s-j+2)}{s\sqrt{p}}, & j = 2, \ldots, s+1 \\ 0, & \text{otherwise.} \end{cases}$$

It corresponds to the DCT coefficients of a (rescaled) sinc^2 function. θ_2^* is the signal that is $(s+1)$-sparse with uniform-value non-zero coefficients and the location of the non-zeros being random. In particular, θ_2^* belongs to the packing set $\mathcal{H}_{\alpha_s}$ with DCT basis. θ_3^* is:

$$(\theta_3^*)_j = \begin{cases} \frac{1}{\sqrt{p}}, & j = 1 \\ \frac{2^{k/2}}{\sqrt{p}}, & j = 2^k + 1, k = 0, \ldots, \min(s, \log_2 p) - 1 \\ 0, & \text{otherwise.} \end{cases}$$

When $s \geq \log_2 p$, θ_3^* corresponds to the DWT coefficients of signal $[1, 0, \ldots, 0]^\top$. θ_4^* is the signal in the packing set $\mathcal{H}_{\alpha_s}$ with DWT basis, where each signal has s randomly located uniform-value coefficients.

B.2 Calculation of λ_k

B.2.1 Discrete cosine transform (DCT)

For a DCT transform matrix D^{DCT},

$$D_{i,j}^{\mathrm{DCT}} = \sqrt{\frac{2}{p}} \cos\left(\frac{(2i-1)(j-1)\pi}{2p}\right), i, j = 1, \ldots, p. \tag{42}$$

For $1 \leq k \leq s$, we can upper bound $\lambda_k(\bar{D})$ and $\lambda'(\bar{D})$ using

$$\lambda_k\left(\bar{D}^{\mathrm{DCT}}\right) \leq \sqrt{\frac{2}{p}} \sum_{j=2}^{k+1} 1 = \frac{\sqrt{2}k}{\sqrt{p}}.$$

The upper bound for λ_k is tight when p is large. To see this, fix $i = 1$, so we get a lower bound for the two quantities:

$$\lambda_k\left(\bar{D}^{\mathrm{DCT}}\right) \geq \sqrt{\frac{2}{p}} \sum_{j=2}^{k+1} \left|\cos\left(\frac{(j-1)\pi}{2p}\right)\right| \approx \sqrt{\frac{2}{p}} \sum_{j=2}^{k+1} \cos 0 = \frac{\sqrt{2}k}{\sqrt{p}},$$

when $p \gg k$.

B.2.2 Discrete Hadamard transform (DHT)

The discrete Hadamard transform matrix $D_{(m)}^{\mathrm{DHT}} \in \mathbb{R}^{2^m \times 2^m}$ has the form

$$\left(D_{(m)}^{\mathrm{DHT}}\right)_{i,j} = \frac{1}{2^{\frac{m}{2}}}(-1)^{i \cdot j}. \tag{43}$$

Thus, because $p = 2^m$, we have

$$\lambda_k\left(\bar{D}_{(m)}^{\mathrm{DHT}}\right) = \frac{1}{2^{\frac{m}{2}}}\sum_{i=2}^{k+1} 1 = \frac{k}{2^{\frac{m}{2}}} = \frac{k}{\sqrt{p}}.$$

B.2.3 Discrete Haar wavelet transform (DWT)

For a Haar wavelet transform matrix of dimension 2^m by 2^m, the non-zeros have magnitudes in the set $\{2^{-m/2}, 2^{-(m-1)/2}, \ldots, 2^{-1/2}\}$. Thus λ_k will be a sum of a geometric sequence (let $m' = \min(k, m-1)$):

$$\lambda_k(\bar{D}_{(m)}^{\mathrm{DWT}}) = \frac{1}{2^{\frac{m}{2}}}\sum_{i=1}^{m'} 2^{\frac{m-i}{2}} = \frac{1}{2^{\frac{m}{2}}}\left(\frac{2^{\frac{m-1}{2}}(1-(\frac{1}{\sqrt{2}})^{m'})}{1-\frac{1}{\sqrt{2}}}\right) = \frac{1-2^{-\frac{m'}{2}}}{\sqrt{2}-1} \to \frac{1}{\sqrt{2}-1},$$

as m' goes to infinity.

B.3 Comparison of compressed sensing to the downsampling method

B.3.1 Setup

Assume f^* is sparse under the DWT basis, and θ^* has exactly $s+1$ non-zero elements. Again, s is the number of non-zeros in the non-DC coefficients. Note that in order to make the downsampling scheme to work well, we need the signal to bear a certain amount of smoothness. We separate the non-DC coefficients into coarse scales ($\theta_i^*, i = 2, \ldots, p/\kappa$, low-frequency coefficients) and fine scales ($\theta_i^*, i = p/\kappa + 1, \ldots, p$, high-frequency coefficients). The downsampling process will discard all of the information in the fine scales, but directly sense the coarse-scale coefficients. Let s' be the number of non-zeros in coarse scales , λ be our sparse-localization constant ($\approx \frac{1}{\sqrt{2}-1}$ and does not depend on s for DWT, so we omit the subscript here). We further assume that all non-zero non-DC coefficients have the same magnitude so that the signal energy in the coarse scales is non-trivial and can be reflected by s'. Specifically, we use signals from the packing set $\mathcal{H}_{s,\alpha_s}$ where $\alpha_s = \frac{1}{p\lambda}$.

Let A^{DS} and A^{CS} be the sensing matrices using downsampling and compressed sensing methods defined in Eq. (11) and Eq. (3) (4), respectively; $y^{\mathrm{DS}} \sim \mathrm{Poisson}(A^{\mathrm{DS}} f^*)$ and $y^{\mathrm{CS}} \sim \mathrm{Poisson}(A^{\mathrm{CS}} f^*)$ be the corresponding observations. Let $\widehat{\theta}^{\mathrm{DS}}$ and $\widehat{\theta}^{\mathrm{CS}}$ be the reconstructions using downsampling and compressed sensing methods, respectively. For compressed sensing, we use the penalized maximum likelihood estimator which has been introduced in Section 2.3.

Similar to our main results, we evaluate the MSE which is defined as

$$R^{\mathrm{DS}} \triangleq \mathbb{E}\left[\|\widehat{\theta}^{\mathrm{DS}} - \theta^*\|_2^2\right],$$
$$R^{\mathrm{CS}} \triangleq \mathbb{E}\left[\|\widehat{\theta}^{\mathrm{CS}} - \theta^*\|_2^2\right].$$

B.3.2 Upper bound of downsampling

In the following lemma we provide an upper bound of the downsampling method for the set of signals we use in the simulations.

Lemma B.1 (Upper bound of the downsampling method)**.** *Let D be the DWT basis. For all $\theta^* \in \mathcal{H}_{s,\alpha_s}$ with $\alpha_s = \frac{1}{p\lambda}$, if there are $s' \leq s$ non-zeros in the coarse scales, and we use the sensing/reconstruction method defined in Eq.* (11) *and* (12)*, then*

$$R^{\mathrm{DS}} \leq \frac{2s'}{Tp\kappa} + \frac{s - s'}{\lambda^2 p^2}.$$

The proof is provided in Appendix A.3.

For the same set of signal $\mathcal{H}_{s,\alpha_s}$, assume we are in the low-intensity region, where the lower bound gives

$$R^{\mathrm{CS}} \geq \frac{s}{\lambda^2 p^2}. \tag{44}$$

Thus the difference is at least

$$R^{\mathrm{CS}} - R^{\mathrm{DS}} \geq \frac{s'}{\lambda^2 p^2} - \frac{2s'}{Tp\kappa},$$

which is larger than zero (downsampling performs better) when $T/p > c/\kappa$ for some constant c. Also, note that this difference is proportional to s', which means when the signal energy in the coarse scales is larger, the performance difference is also larger.

References

[1] M. Raginsky, R. Willett, Z. Harmany, and R. Marcia. Compressed sensing performance bounds under Poisson noise. *IEEE Transactions on Signal Processing*, 58(8):3990–4002, 2010. arXiv:0910.5146.

[2] E. J. Candès and T. Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Transactions on Information Theory*, 52(12):5406–5425, 2006.

[3] D. L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006.

[4] M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly, and R. G. Baraniuk. Single pixel imaging via compressive sampling. *IEEE Signal Processing Magazine*, 25(2):83–91, 2008.

[5] V. Studer, J. Bobin, M. Chahid, H. S. Mousavi, E. Candès, and M. Dahan. Compressive fluorescence microscopy for biological and hyperspectral imaging. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 109(26), 2012.

[6] D. Snyder. *Random Point Processes*. Wiley-Interscience, New York, NY, 1975.

[7] Z. T. Harmany, J. Mueller, Q. Brown, N. Ramanujam, and R. Willett. Tissue quantification in photon-limited microendoscopy. In *Proc. SPIE Optics and Photonics*, 2011.

[8] J. Bobin, J.-L. Starck, J. Fadili, Y. Moudden, and D. L. Donoho. Morphological component analysis: An adaptive thresholding strategy. *IEEE Transactions on Image Processing*, 16(11):2675–2681, 2007.

[9] J. Bobin, J.-L. Starck, J. Fadili, and Y. Moudden. Sparsity and morphological diversity in blind source separation. *IEEE Transactions on Signal Processing*, 16(11):2662–2674, 2007.

[10] C. Estan and G. Varghese. New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice. *ACM Transactions on Computer Systems*, 21(3):270–313, 2003.

[11] Y. Lu, A. Montanari, B. Prabhakar, S. Dharmapurikar, and A. Kabbani. Counter Braids: A novel counter architecture for per-flow measurement. In *Proc. ACM SIGMETRICS*, 2008.

[12] M. Raginsky, S. Jafarpour, R. Willett, and R. Calderbank. Fishing in Poisson streams: Focusing on the whales, ignoring the minnows. In *Proc. Forty-Fourth Conference on Information Sciences and Systems*, 2010. arXiv:1003.2836.

[13] M. Raginsky, S. Jafarpour, Z. Harmany, R. Marcia, R. Willett, and R. Calderbank. Performance bounds for expander-based compressed sensing in Poisson noise. *IEEE Transactions on Signal Processing*, 59(9), 2011. arXiv:1007.2377.

[14] L. Sansonnet. Wavelet thresholding estimation in a Poissonian interactions model with application to genomic data. *Scandinavian Journal of Statistics*, 2013. doi: 10.1111/sjos.12009.

[15] J. M. Boone, E. M. Geraghty, J A. Seibert, and S. L. Wootton-Gorges. Dose reduction in pediatric CT: A rational approach 1. *Radiology*, 228(2):352–360, 2003.

[16] J.-M. Xu, A. Bhargava, R. Nowak, and X. Zhu. Socioscope: Spatio-temporal signal recovery from social media. *Machine Learning and Knowledge Discovery in Databases*, 7524, 2012.

[17] D. W. Osgood. Poisson-based regression analysis of aggregate crime rates. *Journal of Quantitative Criminology*, 16(1), 2000.

[18] E. Candès and T. Tao. Decoding by linear programming. *IEEE Transactions on Information Theory*, 51(12):4203–4215, December 2005.

[19] G. Raskutti, M.J. Wainwright, and B. Yu. Minimax rates of estimation for high-dimensional linear regression over ℓ_q-balls. *IEEE Transactions on Information Theory*, 57(10):6976–6994, 2011.

[20] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.

[21] P. Reynaud-Bouret. Adaptive estimation of the intensity of inhomogeneous Poisson processes via concentration inequalities. *Probability Theory and Related Fields*, 126(1):103–153, 2003.

[22] T. S. Han and S. Verdu. Generalizing the Fano inequality. *IEEE Transactions on Information Theory*, 40:1247–1251, 1994.

[23] I. A. Ibragimov and R. Z. Has'minskii. *Statistical Estimation: Asymptotic Theory*. Springer-Verlag, New York, 1981.

[24] Z.T. Harmany, R.F. Marcia, and R.M. Willett. This is SPIRAL-TAP: Sparse Poisson Intensity Reconstruction ALgorithms - Theory and Practice. *IEEE Transactions on Image Processing*, 21(3):1084–1096, 2012.

[25] J. N. Laska and R. G. Baraniuk. Regime change: Bit-depth versus measurement-rate in compressive sensing. *IEEE Transactions on Signal Processing*, 60(7):3496–3505, 2012.

[26] Y. Plan and R. Vershynin. Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach. *IEEE Transactions on Information Theory*, 59(1):482–494, Jan 2013.

[27] S. A. Van de Geer. High-dimensional generalized linear models and the LASSO. *The Annals of Statistics*, 36(2):614–645, 2008.

[28] I. Rish and G. Grabarnik. Sparse signal recovery with exponential-family noise. In *Compressed Sensing & Sparse Filtering*, pages 77–93. Springer, 2014.

[29] S. M. Kakade, O. Shamir, K. Sridharan, and A. Tewari. Learning exponential families in high-dimensions: Strong convexity and sparsity. *arXiv preprint arXiv:0911.0054*, 2009.

[30] I. Rish and G. Grabarnik. Sparse signal recovery with exponential-family noise. In *Communication, Control, and Computing, 2009. Allerton 2009. 47th Annual Allerton Conference on*, pages 60–66. IEEE, 2009.

[31] J. Jia, K. Rohe, and B. Yu. The LASSO under Poisson-like heteroscedasticity, 2010.

[32] Y. Yang and A. Barron. Information-theoretic determination of minimax rates of convergence. *Annals of Statistics*, 27(5):1564–1599, 1999.

[33] J. Q. Li and A. R. Barron. Mixture density estimation. *In Advances in Neural Information Processing Systems*, 12, 2000.

[34] E. D. Kolaczyk and R. D. Nowak. Multiscale likelihood analysis and complexity penalized estimation. *Annals of Statistics*, pages 500–527, 2004.